AD-A137 343 INTELLIGENCE FUSION MODELING A PROPOSED APPROACH(U) 1/1
MITRE CORP MCLEAN VA MITRE C3I DIV H J ANTONISSE
16 SEP 83 WP-83W00379 F19628-82-C-0001

UNCLASSIFIED F/G 15/4 NL

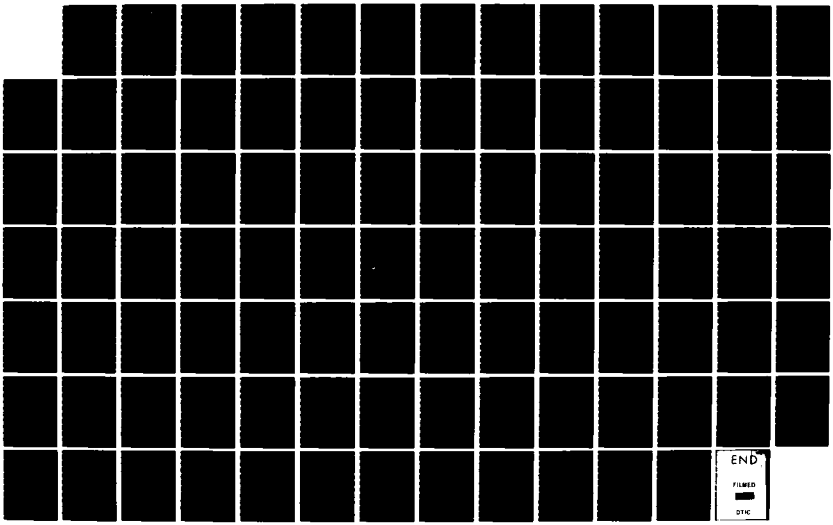

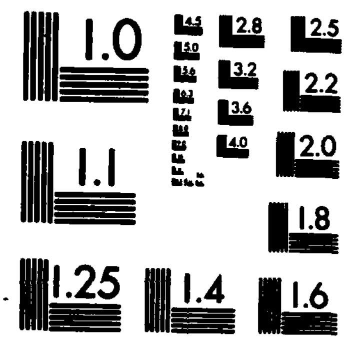

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

(1)

MITRE

WORKING PAPER

The MITRE Corporation
MITRE C^3I Division
Washington C^3I Operations
1820 Dolley Madison Boulevard
McLean, Virginia 22102

WP- 83W00379

No.	Vol.	Series	Rev.	Supp.	Corr.

CONTROLLED DISTRIBUTION

Subject: Intelligence Fusion Modeling - A Proposed Approach

To: Manfred Gale

From: H. James Antonisse

Contract No.: F19628-82-C-0001

Dept.: W74

Sponsor: AMMO

Date: 16 September 1983

Project No.: 8608A

DTIC ELECTE JAN 30 1984 S

Approved for MITRE Distribution: Manfred Gale

ABSTRACT: The U.S. Army is in urgent need of improved methods of modeling Intelligence and Electronic Warfare (IEW) in its battlefield simulations. We present a detailed methodology for modeling the all source fusion part of IEW. The approach draws on Artificial Intelligence techniques that have been used to successfully model human expertise in a variety of areas. A rule-based architecture with highly structured knowledge and data representations is developed. It will automatically correlate and integrate reports from different kinds of intelligence sources, respond to intelligence requests such as Primary Information Requirements and other Information Requirements (PIR/IR), keep requesting agencies appraised of changes in the perception of the battlefield, and justify its actions and answers.

The development of the fusion methodology is in support of the Army Model Improvement Program (AMIP) through the AMIP Management Office (AMMO). The present work draws heavily on previous work sponsored by AMMO.

THIS INFORMAL PAPER PRESENTS TENTATIVE INFORMATION FOR LIMITED DISTRIBUTION.

ABSTRACT

The U.S. Army is in urgent need of improved methods of modeling Intelligence and Electronic Warfare (IEW) in its battlefield simulations. We present a detailed methodology for modeling the all source fusion part of IEW. The approach draws on Artificial Intelligence techniques that have been used to successfully model human expertise in a variety of areas. A rule-based architecture with highly structured knowledge and data representations is developed. It will automatically correlate and integrate reports from different kinds of intelligence sources, respond to intelligence requests such as Primary Information Requirements and other Information Requirements (PIR/IR), keep requesting agencies appraised of changes in the perception of the battlefield, and justify its actions and answers.

The development of the fusion methodology is in support of the Army Model Improvement Program (AMIP) through the AMIP Management Office (AMMO). The present work draws heavily on previous work sponsored by AMMO.

INTELLIGENCE FUSION MODELING - A PROPOSED APPROACH

TABLE OF CONTENTS

TABLE OF CONTENTS

(CONCLUDED)

LIST OF ILLUSTRATIONS

LIST OF TABLES

1.0 INTRODUCTION

The modeling of the activities of an intelligence fusion center is a difficult problem. In a joint Army Research Institute/Intelligence and Security Command (ARI/INSCOM) briefing on a proposed Intelligent Analyst's Manual,[1] the Intelligence Analyst's job was described as follows:

"What is the (Intelligence Analyst's) job like?

<u>Answer:</u> It is very complex.

- There are many clients with many information needs.
- Enemy threat is composed of complex entities.
- The foreign environment is complex and uncertain.
- Collection systems use complex technology.
- The intelligence system organization is complex.
- Analytic processes are complex.
- Communication processes used to acquire and transmit information are complex."

Similarly difficult problem environments have been made tractable through the use of knowledge-based techniques developed by artificial intelligence researchers. This paper describes the application of these techniques in the modeling of an intelligence fusion center. This modeling approach is used to both build up a representation of the battlefield and to respond to information requests that relate to fighting the battle. It can also be used, in principle, for the processing of actual intelligence reports from the battle area, and will be evaluated as a prototype for future system development and procurement.

1.1 Purpose

The purpose of this document is to present a methodology for modeling All-Source Intelligence Fusion based on Artificial Intelligence (AI) techniques. All-Source Intelligence Fusion is the process of correlating, analyzing, and integrating information originating from the diverse collection resources that support the modern battle force. This effort is in support of the force commander who needs to "see" the battlefield, determine enemy intentions, project the impact of the environment on the battlefield, evaluate the progress of the battle, and support the battle in order to operate effectively.[2] This paper describes a system that models the integration of multiple information sources in order to develop a picture of the battle as well as to support a force commander's specific intelligence needs. As such, it fills an urgent need in U.S. Army modeling and simulation by pointing the way towards the embedding of models of intelligence analysis in simulation environments.

1.2 Background

Some aspects of battlefield behavior are exceedingly difficult to model. Areas such as Command and Control (C^2) and Intelligence and Electronic Warfare (IEW) are not well-understood and are not easily amenable to the numeric computation techniques that have traditionally been used to build models and simulations.[3] Part of the difficulty has been that commanders and intelligence analysts in the field rely on conceptual problem-solving abilities rather than the numerical problem-solving typically done by computer.

Artificial Intelligence (AI) is a branch of Computer Science that has focused on symbolic computation techniques for reasoning in complex problem domains and has, over the past 25 years, given computers the capability to do many kinds of conceptual problem-solving. AI techniques are most appropriate when the data for solving the problem is incomplete, unreliable, or changing with time, when the knowledge about the domain is uncertain, and when the search space of solutions is very large.[4] In short, AI techniques work best in environments that closely resemble the information profile of the modern battlefield.

1.3 Scope

The scope of this document is bounded by the Functional Area Representation Objectives[5] (FAROs) for IEW and the IEW Model Requirements Definition Document[6] (MRDD), which specify, respectively, the functional contribution of IEW to the battle force as a whole and the capabilities of and environmental effects on the IEW component of the battle force. In this paper we are interested in the intelligence fusion function. The intelligence fusion aspects of IEW are housed, in the real world, in the All Source Production Sections (ASPS) of the corps and the division G2's staff. In the FAROs this corresponds to the Collection Mission Management, Fusion Management, and Fusion aspects of the Corps and Division Evaluation Model (CORDIVEM). In the MRDD used for the IEW Functional Area model, those functions are further divided and distributed among the processes of Primary Information Requirements/Information Requirements (PIR/IR)* Decomposition, Collection Tasking, Single Source Correlation, and Fusion.

Figure 1** shows the IEW elements of a Corps and Division level model of a battle force as found in the CORDIVEM FAROs. The entire model of the force would include similar decompositions for Air Defense Artillery, Fire Support, Combat Service Support, and Maneuver, with Force Control as the integrating element. The IEW decomposition found in the MRDD, Section 3, is a refinement of the IEW elements found in the CORDIVEM FARO. The fusion methodology described in this paper will provide a way of modeling the fusion node for both the IEW Functional Area Model and the IEW portion of the CORDIVEM.

1.4 Report Outline

The paper is presented in six parts. This introduction is followed by a section that defines the problem from three perspectives: the requirements of the U.S. Army modeling community for modeling all-source intelligence analysis, the computational characteristics of the modeling problem, and the special considerations that apply in the modeling domain.

The third section is the heart of the paper. It begins with a short overview of the AI techniques employed, then describes the proposed approach for modeling the intelligence fusion process, and concludes with a detailed discussion of each major component of the technique as applied to fusion modeling.

* PIR/IR were formerly known as Essential Elements of Information/Other Information Requirements (EEI/OIR).

** Figure 1 shows long range reconnaissance patrols (LLRPs) and remotely monitored sensors (REMS's) to be organic to the division. Actually they are only organic to the corps and brigade, respectively.

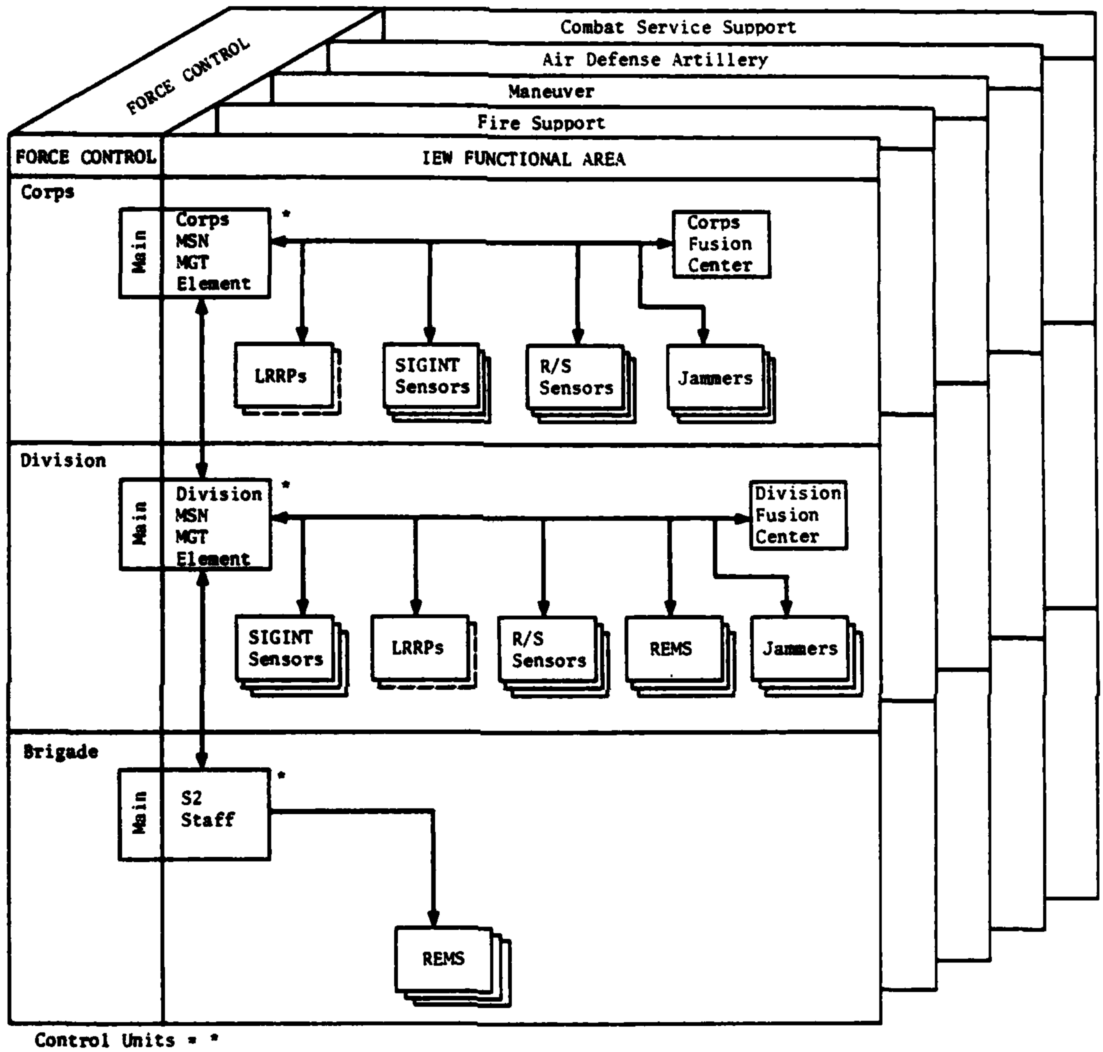

FIGURE 1
I/EW FUNCTIONAL AREA

The fourth section contains an example derived from U.S. Army intelligence doctrine and developed to illustrate the architectural ideas of Section 3.0.

The fifth section briefly discusses some issues that arise in the application of the AI technology to this task, the certification of the knowledge bases in the system, and some hardware implications.

The last section summarizes the characteristics of the proposed methodology and reviews the requirements set forth in Section 2 in light of these characteristics.

2.0 PROBLEM DEFINITION

The intelligence fusion modeling problem is defined by the needs of the user community of Army analysts, the computational characteristics of the problem domain, and the additional considerations due to the nature of the modeling and simulation endeavor. The details of the problem within these areas are described in the following sections.

2.1 Intelligence Fusion Requirements

The specific needs of the Army modeling community for effectively simulating IEW at the Corps and Division level have been distilled in the IEW FAROs and the MRDD. These documents were prepared by MITRE and reviewed by the TRADOC community. They are in support of the Army Model Improvement Program (AMIP) and were produced at the behest of the AMIP Management Office (AMMO). They describe, in particular, the required functional and operational behavior of the intelligence fusion component in the AMIP models. These requirements for modeling fusion are discussed below.

2.1.1 IEW Functional Area Representation Objectives*

In this subsection the functional representation objectives for fusion, i.e. the fusion capabilities essential to IEW, are detailed. Intelligence fusion for corps and division is currently done in the ASPS at both echelons. The ASPS's are referred to here as the corps and division fusion centers. Since the corps and division fusion centers are functionally identical, only the corps description will be presented.

2.1.1.1 Functional Representation of the Corps Fusion Center. The Corps fusion center uses sensor reports of all types along with terrain and weather data to determine enemy location, strength, and intent. The center's own staff and computer databases do detailed correlation and aggregation of the reported data. The center is highly vulnerable to computer damage and the staff involved is very highly skilled and difficult to replace if wounded or

* Much of Section 2.1.1 is taken directly from the FAROs for the CORDIVEM, Appendix IV, IEW Functional Area Representation Objectives.

killed through enemy action. Should the corps fusion center be degraded more than 50% for more than 8 hours, the Echelon Above Corps (EAC) fusion center will establish direct channels with the division fusion centers. The fusion process is an on-going one, officially begun when the corps commander specifies his PIR,* and continued by the commander/staff for information as the battle progresses. During the fusion process, the answering of one question may require the generation of another, thus creating an internal form of tasking as well as external PIR/IR.

The conclusions derived from the multi-source analysis are reported to the corps Collection Management and Dissemination Section (CMDS) for dissemination to the force commanders involved (corps, division and/or brigade). Target development information is also reported to the corps CMDS for transmission to the corps Field Artillery Section (FAS) for targeting, although some sensors have direct channels to the FAS that will circumvent the CMDS.

The corps fusion center will not move more than once or twice per day in the normal course of its operation. This depends, however, on the battle situation. It will displace with the corps main command post when that moves and operate at a degraded level during the movement.

2.1.1.2 Representation of Effects of and on the Fusion Center. The FAROs also list the effects that need to be modeled for an adequate simulation. These include both effects of the fusion activity on other elements in the model and the effects of other activities on the fusion center. In order to do this, the IEW FAROs decompose the IEW functional area into Collection Mission Management, Jamming Mission Management, Fusion Management, Collection (of all types), Jamming, Fusion, and

* Actually, PIR/IR are usually determined by the G2 upon receipt of the Commander's Guidance. The Commander's Guidance is a very high level, general description of a mission. It is beyond the scope of this paper to treat the decomposition of a mission into intelligence requirements, although the methods presented are applicable. We treat PIR/IR as given.

Movement.* These capabilities are described in terms of five major categories of effects that are relevant for modeling and in terms of the execution of the capabilities essential to IEW. Since the effects of executing the fusion capability on targets, the environment, and assets are strictly indirect through effects on command and control decisions, they will not be discussed further. The other categories include combat effects on the capability, environmental effects on the capability, situational factors, and effects from other functional areas. These effects objectives, as relevant to fusion modeling, are summarized below.

Combat Effects on the Capability

The control of IEW functions can be directly affected by enemy conventional and unconventional attacks. The enemy can attack the communications links vital to mission and fusion management with either jammers or anti-radiation guided weapons. He can also attack the management/fusion center(s) through direct ground attack, interdiction strikes, and indirect fire. Nuclear and chemical effects will include blast, nuclear fires, Transient Radiation Effects on Electronics (TREE) and Electro-Magnetic Pulse (EMP) (on the communications and computer equipment), radiation or chemical agent induced sickness, and the degradation of operations while in Mission Oriented Protective Posture (MOPP) status. An added vulnerability particular to ASPS is its reliance on computers containing databases and on fusion algorithms to increase the speed of intelligence processing.

Environmental Effects on the Capability

Communications equipment used in the direction of collection, jamming and fusion can suffer from signal attenuation due to range, terrain, and/or atmospheric considerations. Computers used in the fusion process are also subject to voltage fluctuations caused by lightning.

* The IEW MRDD uses a refined decomposition. See Table 1.

Situational Factors

When the corps or division main command posts (CPs) displace, communications with the sensors/jammers is temporarily disrupted or limited. The analytical direction is similarly disrupted by the movement of the fusion center. The disruption is often limited by a "jump" strategy, where part of the center sets up in the new location before the old location is vacated.

Analysis and synthesis of intelligence reports depends both on the actual volume of reports coming into the fusion center, and on the volume of reports that can be processed by the center's assets. Too few reports restrict the validity of the conclusions made, while too many reports will swamp the analysts and delay conclusions and their reporting. The quality and responsiveness of the fusion process depends greatly on the skill and availability of the fusion analysts and their supporting sensor/jammer analytical teams, as well as the quality and completeness of the reports received.

Effects from Other Functional Areas

Decontamination of personnel and equipment requires combat support (NBC Defense Company) and/or combat service support for supply of water, spraying equipment, and protective garments. The supply of fuel to power communications and computer systems is also dependent on the availability of petroleum, oil, and lubricants.

2.1.2 IEW Model Requirements Definition Document*

The IEW MRDD outlines the user requirements for a complete IEW functional area model in the context of CORDIVEM. It recommends an expanded set of IEW capabilities over the IEW FAROs in order to more accurately model IEW element behaviors. It analyzes these capabilities as a set of generic operations as shown in Table 1.

* Most of Section 2.1.2 is taken directly from the MRDD.

TABLE 1

IEW MODEL PROCESSES VS. CORDIVEM IEW FARO CAPABILITIES

IEW	CORDIVEM
• Situation Development and Target Development	
- Intelligence Preparation of the Battlefield (IPB)	
- Collection Management	Collection Mission Management
- Collection Requirements Decomposition	n/a*
- Collection Tasking	Collection Mission Management
- Collection Monitoring	Collection Mission Management
- Collection	
- Processing	
- Single-Source Correlation	n/a
- Multi-Source Analysis (Fusion)	Fusion
- Target Value Analysis (TVA)	n/a
- Post Attack Assessment	n/a
- Dissemination	Communications
• EW Operations	n/a
- EW Mission Planning and Tasking	Jamming Mission Management
- ESM Operations	Collection
- ECM Operations	Jamming
- Imitative Electronic Deception (IED)	n/a
- Jamming	Jamming
- EW Mission Assessment	n/a
• Movement movement	
• Communications	communications

* Not addressed or not applicable to CORDIVEM scope.

Each capability is thought of as embedded in a program module in the IEW model. The paper describes a methodology for implementing a subset of the capabilities. Section 3 describes a program architecture for a fusion module which embodies the PIR/IR Requirements Decomposition, the Collection Tasking, the Single Source Correlation and the Fusion capabilities. We address collection only insofar as it relates to fusion information needs. We will not be developing a methodology for modeling collection assets and behaviors, only collection requirements. The rest of the IEW capabilities are outside of the fusion domain and will be considered to affect fusion only through inputs to the fusion module.

The IEW MRDD document also provides IEW Process Descriptions. These describe the expanded IEW capabilities noted in Table 1 in terms of an Input-Process-Output template. They are meant as guidelines for later modeling of the capabilities, but do not constrain the models to any particular methodologies. Conversely, this paper will propose specific architectural structures for modeling PIR/IR Decomposition, Single Source Correlation, and multiple source Fusion, and a process structure for Collection Tasking. The IEW Process Descriptions for those three areas appear below

2.1.2.1 Collection Requirements Definition*

- Input. The collection requirements definition phase is triggered by requests from the force control elements for answers to the PIR/IR. In the IEW model, the PIR/IR flow down from force control and must be decomposed into recognizable data items which can be gathered by the IEW collection systems. In a similar manner, high value targets (HVT) are received from the EWS and ASPS which are decomposed into collectable data items.[39] The results from the Intelligence Preparation of the Battlefield (IPB) process are received from the Corps G2.

* Same as MRDD, Section 4.1.2.1.

- Process. The PIR/IR are broken down into the critical indicators, and critical indicators into data elements that can be collected. For example, the PIR posed as "Will the enemy attack in the south, and if so, when?" would be decomposed into the critical indicators of an attack for the situation at hand. Since a critical indicator of an attack is the forward displacement of artillery units, the resulting data elements required to answer the PIR would be:
  - Number of artillery units in the southern sector
  - Location of movement, and speed
  - Some criteria for establishing that they are "forward"
  - An estimate of when that state will be achieved
- Output. The result of this process is a list of collection requirements at the data element level keyed to the PIR/IR/HVT which sparked the process. These requirements are given to the collection management element. As the data elements are found, critical indicators can be verified, and the PIR answered.

2.1.2.2 Single-Source Correlation Processing*

- Input. Critical indicators and data items from the PIR/IR/HVT decomposition are received from the CMDS. Formatted tactical intelligence reports are received from sensors of like types. Maintenance and performance histories of sensor systems are maintained by the Technical Control and Analysis Center (TCAE).
- Process. Reported data elements are reviewed for consistency and validity, and are checked against known sensor error characteristics. Data from individual sensors is compared witht that from other sensors tasked in the same area to determine overlaps and/or confirmations. PIR/IR filled at this level are reported to the collection management authority (CMDS).

* Same as MRDD, Section 4.1.4.1

- Output. Fulfilled PIR/IR are sent to the CMDS. Collected and/or corrected data is forwarded to the ASPS for further processing. Satisfied HVT are passed to the Corps Field Artillery Section (FAS) (or the Corps Electronic Warfare Section (EWS) for electronic HVT) for targeting.

2.1.2.3 Fusion*

- Input. Tactical intelligence reports from the single-source correlation phase, weather reports, and critical indicators formed through requirements decomposition effort are sent to the fusion module for multi-source analysis.
- Process. Data elements gathered from various sources are time ordered, overlaid, and cross-correlated. Evaluations are performed to eliminate obvious errors, reliability factoring is done to assess relative confidence levels, and critical indicators are either satisfied or not, depending on the success of the collection effort.** Targeting requirements are also reviewed for those targets not found with single-source methods.
- Output. Satisfied PIR/IR/HVT are assembled for dissemination to the force commander and his staff. Unsatisfied PIR/IR/HVT are evaluated to see if they require further collection, or if they can be removed from the cycle if no longer required.

2.2 Problem Definition from a Computational Perspective

The AMMO IEW requirements and definition documents characterize a computational problem of considerable complexity. In the following sections the information processing requirements implicit in those documents are stated. The information profile of the intelligence fusion task has had very important implications for the fusion software design. In particular, the computational complexity of that task strongly suggests the use of AI methods

* Same as MRDD, Section 4.1.4.2

** In depth analysis is required to support these processes. This paper is mainly devoted to developing an architecture that models this performance.

in the design. The input-processes-output template used in the last section is continued here.

2.2.1 Input

The inputs to a fusion module will be of three types: reports, requests, and system parameters. Reports are communications of intelligence-related events from other modules of the IEW model environment. Requests are intelligence information demands from other modules. System parameters are factors derived from some part of the model which change the fusion module's performance in some indirect way.

2.2.1.1 Reports. The most problematical aspect of the intelligence fusion problem is the data on which its solution depends. However, the nature of the data also suggests the set of software technologies that can be used to formulate a solution. Below we describe the relevant aspects of the reports, and the data on which intelligence is based, and note the constraints they impose on the software solution.

Incomplete Information

The reports from which a picture of the battlefield is to be constructed are not comprehensive. Only a percentage of the events occurring at any time occur where a blue (friendly) sensor is active, and only a percentage of those are noted by the sensor. The intelligence fusion process must therefore employ plausible reasoning to fill the gaps. The intelligence products of the fusion module will be more inference-dependent, that is, they will require more reasoning by assumptions when there are fewer reports to support the constructed picture of the battlefield. Inference-dependent intelligence may be less reliable than products based more on specific reports and known enemy behavior. The base of support for intelligence products must therefore be explicit both to indicate the reliability of the product and to update the product when missing information becomes available.

Unreliable Information

The reports on which intelligence fusion depends may themselves be unreliable. This is most commonly modeled by error elipses for sensors and reliability grades for human intelligence. Therefore the fusion module must do probabilistic as well as plausible reasoning. The fusion process also has to carry out statistical reasoning and pattern recognition as reflected in clustering techniques* to correlate reports from different sources. This effort to improve the reliability of intelligence by intelligent aggregation** and a multi-spectral approach is an important part of this intelligence fusion methodology.

There is another aspect about the unreliability of reports which requires special attention. Enemy diversions will result in reports that may be highly reliable in terms of sensor error profiles but which represent an unreliable estimate of the enemy's true intentions. It must be possible for the fusion module to discover that it has been mislead and correct its actions in the future. This requires that multiple, possibly contradictory hypotheses be maintained during the process and that the explicit base of support for each hypothesis (i.e., an audit trail) be made available.

Time-Varying Data

The modern battlefield is extremely transient. It is characterized by very high activity and mobility of the combatants, making situation assessment difficult and tenuous. Estimates of enemy deployment based on intelligence reports often degrade rapidly in reliability as a function of time. Nonetheless, intelligence fusion will depend on a stream of such perishable reports. This requires the reports to be time-tagged and that some procedure for assessing a report's deteriorating status be available. The delays in

* Clustering techniques: techniques for assessing the similarity of elements in a group of samples in order to collect them into coherent subgroups.

** Clustering in the intelligence domain requires that samples are weighted in a context dependent way. For example when the target location given by a very accurate sensor is supported by (clustered with) a report from a much less accurate sensor, the location given by the second sensor is ignored.

processing the reports into a situation analysis must also be accounted for by requiring the analysis itself to be time-tagged and of deteriorating reliability.

Although knowledge about time is generally under-exploited in AI systems, it is crucial that it be invoked to adequately capture intelligence assessments of expected enemy behavior. In battlefield intelligence, as in many other areas of analysis, a fulfilled expectation is powerful confirming evidence for a hypothesis. An explicit representation for expectations in time provides a powerful tool for modeling human problem solving in the combat intelligence arena.

Implicit Structure in Data

Individual reports may vary widely with respect to the scope of their content. For example, a report from Theater-level intelligence may state "Soviet 3rd Army is moving North along the Hunfeld-Bad Hersfeld corridor", while a moving target indicator (MTI) report may indicate "A platoon sized tank unit is moving North along the Meissenbach-Odensachsen corridor". The report structures are the same, but the unit structures being reported are orders of magnitude apart, as are the sizes of the geographical areas where they are being placed. This requires corresponding structures in the fusion model to process the respective report contents. In the realm of Order of Battle (OB) intelligence, this requires a hierarchical model in the fusion module to distinguish and relate reports dealing with different aspects of the enemy military organization.

2.2.1.2 Requests for Intelligence. Requests for intelligence correspond to the commander's Primary Information Requirements (PIR) and his staff's Information Requirements (IR) in the actual battle environment. PIR are requests such as "Does the enemy intend to attack my south flank in the next 48 hours?" IR are requests for more static information, such as "How many tank battalions face the 3rd Brigade." In contrast to reports, intelligence requests will be treated, for modeling purposes, as clear and unproblematic in content. In the modeling environment, it is possible to specify a concise formal language that will be identical at the output of the IEW module making an intelligence request and at the input of the fusion module. We will not treat the possible garbling of requests, message ambiguities, and other such problems, although they could be modeled using operational parameters (Section 3.2.4).

Requests will usually be time imperative and perishable. Their perishability is similar to report perishability, except that their decay must be monitored. Consider an PIR that asks "Will the enemy strike at the south flank within the next hour." If the system has already determined, recently, that the enemy will/will not attack in that time, the request is quickly answered. If not, the fusion module must proceed to determine the answer from analysis. This may fail to return an answer for lack of information or may return an answer which is too unreliable. When this happens, the fusion module must task the collection module of the IEW model to provide extra information. This requires that a substantial amount of time be invested in answering the PIR. Since the utility of a PIR often deteriorates rapidly, the fusion module must be able to cut off further efforts on it and report partial results if collection is taking too long.

2.2.2 Process

There need to be two modes of processing fusion module inputs, corresponding to the intelligence reports and the intelligence requests. System parameters will in general only affect the processing indirectly. In Section 3, special attention is given to those instances where this does not hold.

Reports

When a report is received, the fusion module must ascertain the reliability of the report, its consistency with previously reported and developed information, its relevance to priority tasks, and must determine a storage and access strategy for the report.

Requests for Intelligence

When an intelligence request is received the fusion module tries to fulfill the request by finding the required information in its databases of known or perceived information. If this fails, it tries to deduce the information using the databases and inferential knowledge about the battlefield. If this also fails, the fusion module generates a request for the collection module to obtain the desired information.

2.2.3 Output

The outputs of the fusion module must be clear, targeted, independent, and qualified. All intelligence products must have a likelihood level attached to them to indicate the estimated reliability of the product as determined by the fusion process. The output is treated here as free of transmission errors and communication failures. This is analogous to the input case, and follows the strategy for developing the fusion methodology independent of other model components and considerations.

Outputs will be addressed to the module requesting the information or, for spontaneous reports, to the agency to whom the information is deemed relevant (See Section 3.2.2.2, System History) because of previous requests to the fusion center. All other considerations of dissemination are outside the scope of the present effort.

Intelligence products are independent in the following important way. The response of the fusion process to a specific request at one time is not necessarily identical to the response to the identical request at another time even if the requests are very close in time. They will be the same, however, if no new information is derived in the interval between the two requests.

2.3 Implications Due to the Modeling Domain

AMMO is responsible for improving the performance of the U.S. Army simulation and modeling effort. The methodology we will describe was developed in the context of AMMO's concern with simulation, interactive wargaming, and training models. The main impact of the modeling considerations on our design was to direct us toward the "rule-based" methods of AI as our primary modeling technique.

2.3.1 Simulation

The fact that the fusion process will be part of a computer simulation imposes special constraints on the system design. These concern time compression, fidelity, and justifiability.

2.3.1.1 Time Compression. The fusion process may have to run considerably faster than real time for effective use in some battlefield simulation models. This requires that any proposed methodology be capable of speed-up, and, ideally, of operating at multiple levels of abstraction.

2.3.1.2 Fidelity. Fidelity is the faithful reproduction of an original. It is the goal of modeling in general. In the context of a knowledge-based computational methodology for fusion, however, this takes on special meaning. The software technology that will be presented for the intelligence fusion task has been used in a number of software systems that have performed better than humans in the task for which they were designed. Since the fusion module is to represent the functions of actual, manned fusion centers in the field, the module should perform no better and no worse than its real-life counterpart. Moreover, if All Source Production sections tend to degrade as a function of the volume of incoming reports or the length of continued operation, the same phenomenon must be reflected in the fusion module's performance. We discuss some non-parametric techniques for attaining such effects in Section 3.3.4, System Parameters.

2.3.1.3 Justification. Justification is the demand that the "lines of reasoning" of the fusion module be accessible and understandable. The results of the fusion process are ultimately targeted for human analytical efforts. In order to both assess system performance and to evaluate tactical and strategic alternatives, the human analyst must have access to the "lines of reasoning" followed by the fusion software. The program trace must therefore be accessible and understandable to the user community.

2.3.2 Interactive Training

The use of the model for training purposes imposes an additional set of requirements on the fusion module in terms of fidelity, expertise, and justifiability. If the fusion module is in support of a personnel training program in battlefield intelligence, the system has to represent the best intelligence techniques available, rather than the best performance an actual fusion center can achieve in the heat of battle. Therefore fidelity is not necessarily valued in a training system. However, justification is even more imperative in the context of training than in simulation. The trainee must be encouraged to question and understand the intelligence process.

2.3.3 War-Gaming

The design decision to use AI techniques is largely based on the fact that human intelligence and problem-solving capabilities are being modeled. It is relatively easy, therefore, to send the sensor stream to a team of human intelligence analysts instead of the fusion module and have them develop answers to PIR/IR. However, the methodology presented here is not a workstation model. That is, it does not claim to emulate an ASPS person for person. It may not be possible to insert specific members of an analyst team into the fusion module for "intra-module" gaming. Therefore, the fusion methodology presented will directly support war gaming at the IEW level, but not, without careful tailoring, at the ASPS. Replacing the module with a team of analysts would require timing adjustments with the other modules, and the formal language requirements for output to other modules would still be in effect.

3.0 METHODOLOGY

We have reviewed the Army requirements for an intelligence fusion model and the computational implications of this requirement. We have also noted the constraints put on any fusion software design by the user community, i.e. the modelers, trainers, and war-gamers who will use the software product. The rest of this paper is devoted to describing a software architecture that can meet the requirements within the constraints described above.

3.1 Knowledge-Based Systems - An Overview

The intelligence fusion methodology is derived from the subfield of AI called knowledge-based systems. In a knowledge-based system, knowledge of the problem domain is deliberately separated from the system's control structure (the method of applying the knowledge to the facts of a case) as well as from specific facts about the current domain. These design characteristics make knowledge based systems easy to modify and capable of explaining their lines of reasoning. To modify them, knowledge in the form of "if-then" rules is added to or deleted from the knowledge base. Lines of reasoning leading to a conclusion are explained by the presentation of the chain of inferences used to derive the conclusion and the facts supporting the inferences.*

The three interacting components of knowledge-based systems are knowledge bases, where domain knowledge is encoded, databases, where the facts known to the system are stored, and the inference engine, which interprets inputs to the databases, extracts facts or applies knowledge to facts to generate results, and dispatches findings as outputs. These components are described in the following paragraphs.

* This technology has been under development for over 20 years and is well understood. There is a large amount of supporting literature on this program architecture [7,8] and a growing familiarity with the techniques of knowledge-bases systems that makes this design path the recommended approach.

3.1.1 Knowledge Bases

Knowledge in knowledge-based systems is most often encapsuled in an if-then rule format. Such a system is also called a rule-based system. An example of a piece of knowledge about Tactical Nuclear Weapons "chunked" in rules is the following:

Tactical Nuclear Weapons (TNW) rule 1:

(Antecedent)

<u>**If**</u> the enemy's SCUD (Soviet tactical rocket) units are within 20 km of the Forward Line of Troops (FLOT)

<u>**and**</u> the forward enemy troop positions have hardened positions with overhead cover

(Consequent)

<u>**Then**</u> this is strongly suggestive evidence that the enemy will employ tactical nuclear weapons (with 80% confidence)

The idea is that the system's knowledge that enemy SCUD units have moved to within 20 km of FLOT, in conjunction with knowledge that the enemy's forward troops are in hardened positions with overhead cover, allows the system to derive that the enemy will probably employ tactical nuclear weapons by a single application of modus ponens.* The confidence is a measure of how much we trust the rule.

Knowledge based systems have been successful not only because of the specialized knowledge they bring to bear on a problem, but also because they are capable of explaining the procedures they use and the results they produce. As was mentioned previously, this explanation capability is achieved by reference to the rules that drive the problem solution. This is done as follows: Suppose a rule-based system for intelligence fusion is trying to establish the antecedent of rule TNW1. It may ask the user "Are enemy SCUDs within 20 km of the Forward Line of Troops (FLOT)?" The user may

* Modus Ponens: Rule of logic that, given A B and A, allows one to conclude B.

reasonably ask "Why?" to which the program can respond "Because I am trying to establish whether the enemy will use tactical nuclear weapons, and by rule TNW1, if the enemy SCUDs are within 20 km of the FLOT and the enemy troops have hardened position, I know that the enemy probably will use tactical nuclear weapons." Similarly, if the system responds that the enemy will use tactical nuclear weapons, the user can ask "How do you know?." A reasonable response by the program would be to display rule TNW1 and indicate the status of the premises. In this way a rule-based system can explain both its actions and its results.*

Rules may sometimes house actions in their consequents. In this way the application of a rule may add or delete facts from the database or even change rules in the knowledge base itself.

3.1.2 Data Bases

In the example above, "enemy SCUD units are within 20 km of the FLOT," was treated as a fact. It was a premise which was known to be true or false by observation rather than by inference. The facts in the database must be represented in the same formal language as the knowledge base so that it can be determined whether the premises of inference rules are established by the facts in the database. It is recommended that the entries of the data base be in canonical form, so that all reports that mean the same thing have the same representation.

3.1.3 The Inference Engine

The knowledge and databases are not self-activating, but are passive representations of procedural and observational knowledge, respectively, about the domain. The inference engine is the software mechanism and control structure that brings these passive structures to bear on domain problems. Input to the system is interpreted by the inference engine in light of the formal language defined by the database and the rules of the knowledge base.

* The interaction is for illustrative purposes. Although it is convenient to have a rule based system be interactive, it is not necessary.

The recursive application of rules to answer a request to the knowledge-based system is called goal-directed backchaining. If, when the antecedents of a production* are matched and the production fires, perhaps changing the database and causing more productions to fire, we have forward-chaining. In either case the chaining depends on the pattern matching, that is, the language commonalities of requests, reports, and rule components. The inference engine will be able to respond to requests couched only in that language.

3.2 The Fusion Methodology

The Fusion Methodology we advocate is essentially a set of knowledge-based systems operating in concert. Figure 2 illustrates the basic scheme. Requests for intelligence arrive from the Force Control or other battle force components in the form of PIR/IR. These are interpreted by the inference engine, which first consults databases and the perceived situation map (Sitmap) to see if it can provide an immediate answer. The Sitmap is a special data base of facts, augmented by inferences that are deemed important or reliable enough to be directly accessible. Failing this, the inference engine begins chaining through its knowledge bases trying to derive the answer to the intelligence request. If this also fails, it determines if it should report its partial results or if it should task collectors for the information it needs to make the assessment. Reports are processed by the inference engine which uses report processing rules to update the Sitmap. Auxiliary data bases are required for facts about enemy weapons characteristics, etc. Results of PIR/IR are output to the original requesting module.

In the next section the required software constructs and processes for the methodology outlined above will be fleshed out. The discussion of the methodology will be organized into four parts: 1) representation of knowledge in the knowledge bases, 2) representation of facts in the database,

---
* production = rule

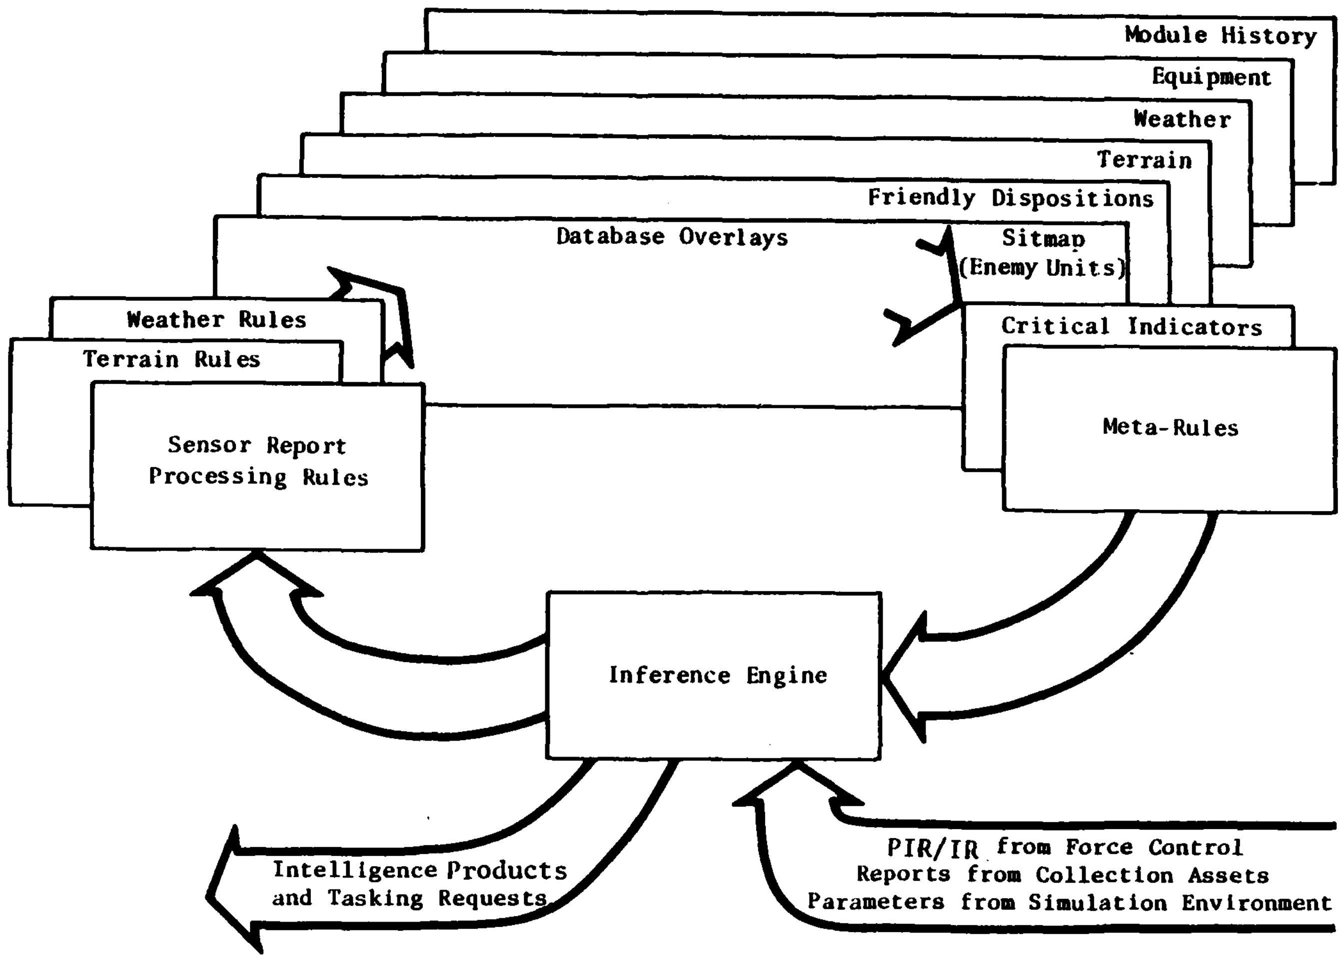

FIGURE 2
A SCHEMATIC OF THE FUSION METHODOLOGY

3) the inference engine control regimen, and 4) system parameters. Each of these parts will include a description of the software structures for implementing the part of the methodology under discussion, and small examples which will illustrate how the design is to operate.

3.2.1 Representation of Knowledge

There are three important features the knowledge representation should have to conveniently model intelligence fusion. Knowledge in the intelligence domain should be hierarchically structured, be semantically factored into roughly independent sub-domains, and be adaptive and self-referencing. The rule-based methodology for knowledge representation described in this section has the desired features. The technique of representation that we advocate is a semantically factored frame-based rule representation with a meta-rule component.

Conceptually, the intelligence fusion process lends itself to partitioning into relatively independent sub-domains. For instance, weather conditions can be reasoned about without detailed analysis of the order-of-battle. The semantic partitioning of the problem results in increased overall efficiency and speed of processing, and makes the fusion process conceptually clearer.

Although it is a sound organizational and implementation strategy, the partitioning of the domain seems to violate the interdependence of information on which successful intelligence fusion depends. An example will show that this is not the case. The partitioning does not compromise the fundamental interdependence of the knowledge areas reflected in the antecedents of the rules. A rule may state:

Critical Indicators (CI) rule 1:

<u>If</u> blue has local air superiority
<u>and</u> there is continued bad weather
<u>Then</u> increase the probability of red attacks
(with 70% confidence)

This rule would be found in the Critical Indicators rule base, since it concludes something about enemy behavior, but it contains statements about weather conditions and friendly force dispositions in its antecedent part. In this way information dependencies are explicitly represented.

The semantic decomposition of the intelligence fusion problem we will follow appears below.

- Report Processing Rules: Procedures for inferring enemy battle units from intelligence reports
- Terrain Rules: The effects of salient topographical features on model elements
- Weather Rules: The effects of weather on model elements
- Critical Indicators: A rule-base of a priori enemy behaviors and indicators of enemy intentions
- Meta-Rules: A set of rules to reason about the system's control behavior, PIR deterioration, tasking, and adaptation

Each knowledge base is described more fully and with examples in the subsequent text. Recommended sources for the knowledge bases are presented in the appendix.

3.2.1.1 Report Processing Rules. In recent years MITRE has developed a knowledge-based system for sensor report fusion known as ANALYST[9]. ANALYST is a forward-chaining production rule program that processes sensor reports onto a situation map of the battlefield.* We incorporate the ANALYST architecture as the report processing component of our fusion methodology. Although another system could be used or a new component developed, employing ANALYST, allows us to draw on a well-established, working technology, and simplifies our fusion problem. ANALYST was implemented as a number of small knowledge-based systems operating together. We present the knowledge-bases as subsets of our Report Processing rule base.

* ANALYST may develop a back-chaining component.

The knowledge base is actually a set of six rule subsets segmented as follows:

- Cluster rules - which gather sensor reports of identical types and similar locations into activity clusters
- Pattern rules - which infer the existence of military units (entities) from patterns of clusters
- Refinement rules - which refine unit attributes from tactical or terrain knowledge
- Merge rules - which determine when to merge two or more inferred units into a single unit with more refined attributes
- Reinforcing rules - which reinforce the inferred existence of enemy units from stray clusters of activity
- Purge rules - which purge hypothesized units from the situation map

A sample from each of the rule sets appears in Table 2.

3.2.1.2Terrain Rules. The terrain knowledge-base holds rules pertaining to salient topographical and environmental features and their effect on model elements. In actual intelligence operations terrain is often a driving consideration in the analysis. In modeling environments, however, this has

often been neglected because of weak or under-exploited representation techniques. We address the terrain representation issue in more detail in the discussion of the databases, below. The terrain representation described there will make possible the efficient application of terrain rules such as:

Terrain Rule 1:

If the sector is swamp

Then it cannot traffic wheeled or heavy vehicles.
(with 90% confidence)

TABLE 2
EXAMPLE REPORT PROCESSING RULES

Cluster Rule:	IF	the received report is COMINT, and the band is VHF, and there are COMINT clusters within 1 KM of the report with the same frequency
	THEN	add the report to the nearest cluster, and reaverage the cluster's location, and repost the cluster to the situation map at the new location
Pattern rule:	IF	there exists a COMINT pattern of at least 3 clusters, and each cluster is composed of at least 2 reports, and at least one of the clusters is in the HF band, and the spread between the maximum report count of all the clusters and the average report count is greater than or equal to 3
	THEN	post a tank battalion COP to the situation map at the center of mass of the COMINT pattern and update the entity statistics
Merge Rule:	IF	an entity exists of known type, and another entity exists of unknown type and the sizes of the two are equal and the two entities are within 1 KM of each other,
	THEN	merge the attributes of the second entity with those of the first, and delete both old entities from the situation map, and post the new entity to the situation map
Refinement Rule:	IF	an entity is of type arty, and the FLOT-Distance of the entity is less than 5 KM,
	THEN	change the entity type to unknown, and repost the entity to the situation map
Reinforcing Rule:	IF	the unused cluster is of type ELINT, and its report count is at least 2, and there is an entity with 1 KM whose type is ADA
	THEN	delete the cluster, and update the last-update time of the entity
Purge Rule:	IF	an entity has a last-update time greater than the purge-time, and the enitty is stationery, and the entity is not a motor transport company,
	THEN	delete the entity from the situation-map

3.2.1.3 Weather Rules. The weather knowledge-base is similar to the terrain knowledge base in that it holds rules pertaining to the impact of the environment on model elements. An example of a weather rule would be:

Weather Rule 1:

If it is raining in a sector

Then MTI reports in the sector are degraded 25% (with 99% confidence)

3.2.1.4 Critical Indicators. In order to successfully analyze and predict enemy activity it is necessary to understand enemy doctrine and order-of-battle. This means that rules concerning enemy behavior must be available. Much of this part of battlefield analysis is embedded in the clustering and pattern rules, but ultimately the statistical and structural aspects of clusters and pattern detection need to be separated. Once they are separated, expertise from the respective areas of pattern recognition and enemy order-of-battle and doctrine may be brought to bear. An example of an order-of-battle rule is:

CI rule 2 (confidence - 85%):

If the unit is a battalion Command Observation Post (COP)

Then there are probably two companies approximately 3-5 KM towards the FLOT and one company within a 2 KM radius of the COP.

Additionally there will be a number of sensitive indicators of enemy activity that give early warning of his intentions. The analysis of enemy intentions amounts, operationally, to the specification of a set of expected enemy behaviors. In order to deal with expected activities, some facility for making tests at future times is needed. An AI technique for achieving this is often called "posting demons." The demons are functions that "wake up," i.e. are called, when an input pattern is matched (such as an expectation being fulfilled or an alarm going off). A demon causes a special action to be taken. The critical indicators knowledge base will consist of general demons that invoke special actions in special circumstances, and specialized demons that are posted by the consequent part of the rules fired in other knowledge bases. An example of a rule of this type is:

CI rule 3 (confidence - 90%):

<u>If</u> there is movement of additional troops towards the front

<u>or</u> increased traffic toward present positions

<u>or</u> new units have been identified in the combat zone

<u>or</u> additional CPs and supply and evacuation installations have been reported

<u>Then</u> set a demon to determine enemy artillery positions in two hours.

3.2.1.5 <u>Meta-Rules</u>. Meta-rules are rules that treat other rules as data. Employing meta-rules can extend the power of a system by giving it a learning capability,[10,11] providing it with an explicit representation for control knowledge,[12] facilitating abstract rule compilation,[13,14] and making it possible for the system to reason about tasking issues.

It will be necessary for the fusion module to have a limited learning capability for fusion to be correctly modeled. This can be achieved by attaching confidences (a percentage in this case) to each rule and incorporating such as:

Meta-rule 1 (confidence - 99%):

<u>If</u> a rule results in a faulty inference

<u>Then</u> (recursively) decrement the confidence of the rule by 25% of its contribution to the inference.

and Meta-rule 2 (confidence - 99%)

<u>If</u> the confidence of an inference rule falls below 15%

<u>Then</u> delete and replace it with a new rule.

When fired, such rules assign blame to the rules immediately leading to a faulty conclusion, as well as the rules that have led to the firing of those rules (see Section 3.2.2.2, System History). Disfunctional rules are eventually replaced by new, hypothesized, but untested rules. When such meta-rules are part of a system that processes large amounts of information and that has the means of evaluating its behavior (i.e adjusting the confidences of its rules) and generating new rules, the systems behavior adapts to its informational environment.[15]

Control issues, such as when to forward-chain and which antecedents to backchain can be kept under explicit program control through a super-level rule-based system that decides when and how the standard rule-based system is executed. Two examples of control rules are:

Meta-rule 3 (confidence - 90%):

<u>If</u> several rules can be applied to establish a goal statement (i.e. several rules have the same consequent).

<u>Then</u> first try the one with the most antecedents.

Meta-rule 4 (confidence - 90%):

<u>If</u> the rule to be backchained has several antecedents

<u>Then</u> establish the (recursively) most meritorious[16] antecedents first.

Meta rule 3 essentially recommends that the most specific rules of a conflict set be tried first. Meta rule 4 orders the backchaining to get the most possible information out of every inference. It also implies that the system has a reliability or certainty threshold that could allow it to stop backchaining before an exhaustive search.

There are several rule-compilation techniques for increasing the storage and computational efficiency of rule-based systems (see, for instance, Van Melle[13] for a discussion of decision-tree compilation). Meta-rules may be fashioned to implement a particularly interesting and effective compilation technique that is closely related to learning by induction. Given an evaluation capability, meta-rules may inductively posit new rules as working hypothesis. Those new rules that work tend to survive, those that don't are pruned. Such applications of meta-rules tend to supplant a given rule-set with a small set that has the same or greater cover[11,14] than the original rules. Meta-rule 5 is an example of such a rule.

Meta-rule 5 (confidence - 90%):

<u>If</u> a rule that mentions a (specific or generic) entity is highly successful

<u>Then</u> add a new rule that is identical to it but that mentions the next most general class of the entity.

Without a meta-rule capability, collection tasking occurs as soon as an information search fails. However, the tasking should be short-circuited if the additional information won't help much, is too costly, will take too long to gather, or if the information already gathered has sufficient support for its intended use. Other uses of meta-rules in tasking are to determine the reliability level that is acceptable and the tasking strategy that will provide the most information at the least time and risk. Examples of each follow:

Meta-rule 6 (confidence - 99%):

<u>If</u> a collection task will improve reliability or confidence of the PIR/IR by less than 10%

<u>and</u> all collections will improve reliability or confidence of the PIR/IR by less than 20%

<u>Then</u> don't task and send the PIR/IR

Meta-rule 7 (confidence - 80%):

<u>If</u> an PIR/IR has a probability of .9 and confidence of .75*

<u>Then</u> send the PIR/IR.

* The probability of a fact is the likelihood of its being true. This can be low even if the rules to determine it can always be used with very high confidence. However, the rules themselves may be suspect, so that one cannot be completely confident of their assertion that a fact is highly probable if it is derived with rules in which we lack confidence.

3.2.2 Fact Representation in the Database

The implementation of the database of system facts will be presented in two parts. The first details a generic representation for facts in the database. The second part describes the database in terms of its structure and its contents.

3.2.2.1 Representation of Facts. The success of a fusion module ultimately depends on its ability to derive, store, and exploit facts about the battle situation. We describe a frame-based[17] implementation of facts that will store much of its information implicitly in relational structures.

In the intelligence fusion domain it is advantageous to use a representation that reflects the relationships and structures of the objects of interest — the enemy fighting organization on the battlefield terrain. Frames can provide such a representation. Frames are structures that have characteristics defined by slots. These slots have names and contain values which may be other frames, lists of frames, numbers, etc. The filling of frame slots with other frames allows defaults and inheritance relations to be specified. The slots in a frame can also be used to specify operations to perform that generate values. This technique is called "procedural attachment" and can be used to represent the effect of the geography and terrain on military units and to do spatial reasoning. This power and flexibility of frame-based representations makes them natural for representing the complex objects of our domain.

Figure 3 illustrates a possible representation for military units that embodies these ideas. In the figure, the class of maneuver units is specified by a frame complete with default attribute values for members of the class. Class membership is indicated by an A-Kind-Of relationship, which indicate from where attribute values can be inherited. For instance, if we ask what kind of mobility an armored unit has, the system first looks for the "Mobility" attribute in the armored-unit frame. Failing there, it looks for an A-Kind-Of attribute, and finds "Mobility" in the maneuver-unit frame specified there. In the 2nd-Armored-Brigade frame, no inheritance or default is invoked for "Mobility" since it is already specified as "very high".

MANEUVER UNIT

```
Type            :   infantry
Superior        :
Subordinates    :
HQ-Location     :   ((4 Km behind Front-sector)
                     (Reliability = .3))
Front-sector    :
Enemy-facing    :   if (null Front-sector)
                    then 0
                    else if (null Subordinates)
                    then  (Strength of Enemy in Front-sector)
                    else  ((Enemy-facing Subordinates)
                            (Enemy-2nd-Echelon of Subordinates))
Weapons         :
Mobility        :   high
A-Kind-Of       :   COMBAT UNIT
```

FIGURE 3a
MANEUVER UNIT FRAME

ARMORED-UNIT

Type	:	armored
Weapons	:	M-60
A-Kind-Of	:	MANEUVER-UNIT

FIGURE 3b
ARMORED UNIT FRAME

2nd-ARMORED BRIGADE

Superior	:	3rd-Division
Subordinates	:	(106th-BN, 107th-BN, 108th-BN)
HQ-Location	:	(6025 7956)
Front-sector	:	((6031 7953) (6029 7961))
Weapons	:	M-1
Mobility	:	very high
A-Kind-Of	:	ARMORED-UNIT

FIGURE 3c
2ND-ARMORED-BDE FRAME

106th-BN

Superior	:	2nd Armored-Brigade
Front-sector	:	((6031 7953) (6031 7956))
Strength	:	45

FIGURE 3d
106TH-BN FRAME

FIGURE 3
FRAME REPRESENTATION EXAMPLE

A request for the strength of the enemy facing the 106th Bn would cause the following sequence of actions. The system fails to find an Enemy-Facing slot in the 106-BN frame. It therefore looks for an A-Kind-Of-slot. This also fails, so it looks for a superior slot from which to inherit the requested information. This succeeds, but when Enemy-Facing is searched for in the new frame environment it fails again. However, 2nd-Armored-Brigade is A-Kind-Of Armored Unit, and the system tries to derive Enemy-Facing from the Armored-Unit frame. This does not hold the required information either, but Armored-Unit is A-Kind-Of maneuver unit, and maneuver unit has an Enemy-Facing slot with a procedure in it. The procedure is passed down to 106th-BN and invoked using the FRONT-SECTOR slot local to it. This finally causes the system to reach out into the correct battlefield sector, and possibly recursively into subsectors, to estimate the strength of the enemy forces against the 106th BN.*

Any relevant intelligence will be keyed to friendly forces or territory. We assume for purposes of simplicity that friendly territory is covered by friendly forces' zones of responsibility. It is most convenient, then, to keep a hierarchical data structure for representing the friendly order of battle and to key intelligence of the enemy unit positions to friendly force dispositions. The integration of the battlefield model is completed by keying the friendly disposition of forces to a hierarchial geographic representation.

This point is important and worth elaborating a little further. A frame of the friendly force is created for, say, a Corps. The Corps has responsibility for a front divided into two sectors, each covered by a division.** Each division front is divided into brigade sectors, those into battalion sectors, and those finally into company-sized sectors of the front. Suppose another "Enemy-facing" request is made, but this time to the 2nd-Armored-Brigade. The system accesses its "Enemy facing" slot

* Notice that in climbing the inheritance graph, we needed to know precedence relations among frame slots. Had "superior" taken precedence over "A-Kind-Of", and had "Enemy-Facing 3rd-Division" been previously established as "Soviet 2nd Army", then the incorrect inheritance of "Soviet 2nd Army" for enemy facing 106th BN" would have occurred. Some things, such as "weapons", are, nevertheless, inheritable through "superior."

** Actually, the Corps is represented as having a perimeter, and its component units as having sub-perimeters, but the "front" representation is a convenient simplification.

and eventually finds the procedure that adds together the entries in the "Enemy units facing self" slots of its battalions. The system then accesses the battalion "Enemy Facing" slots and finds it has to determine Enemy-Facing of the subordinate companies. At the company level there exists an attached procedure that accesses the situation map directly or otherwise finds the known or suspected enemy units near to it. Once found, the results from the company level are passed back up to the battalion level, and the battalion level to the brigade, where they are combined.

The same structure can also automatically organize enemy activity into meaningful intelligence summaries. This is illustrated in Figure 4, which sketches the front decomposed as described above. An enemy attack occurs as shown. At each level the situation is assessed relative to the force sizes at that level. Thus a single very heavy attack at the lowest echelon may be viewed as minor enemy activity several echelons up.

The recommended approach is therefore for the fusion module to have its own data structures for terrain and weather analysis coupled with the military organization data structure. These constitute a hierarchical set of registered terrain "images" of decreasing resolution matched to the echelons of the battle force. The resolution at each level is dependent on the least area a unit of the corresponding echelon would cover. The idea is illustrated in Figure 5.* Such structures, derived from the "pyramids"[18] and "quad trees"[19] of the Computer Vision discipline, allow fast local operations to be employed for determining global unit distributions, enemy activities, etc. These data structures greatly simplify deployment analysis.

The same frame structures used for fact representation may be used for representing knowledge by implementing rules as frames with slots such as "Antecedents" and "Consequents." This streamlines the systems implementation. It

* Notice that the OB representation in Figure 5 involves units rather than fronts. We repeat that the front representation of Figure 4 is purely illustrative and should be thought of as a coupling of unit perimeters. The OB representation by unit perimeters rather than by front sectors facilitates the modeling of pockets and islands of force elements. The natural capability of pyramids to represent these islands makes them especially well suited for modeling Airland Battle doctrine.

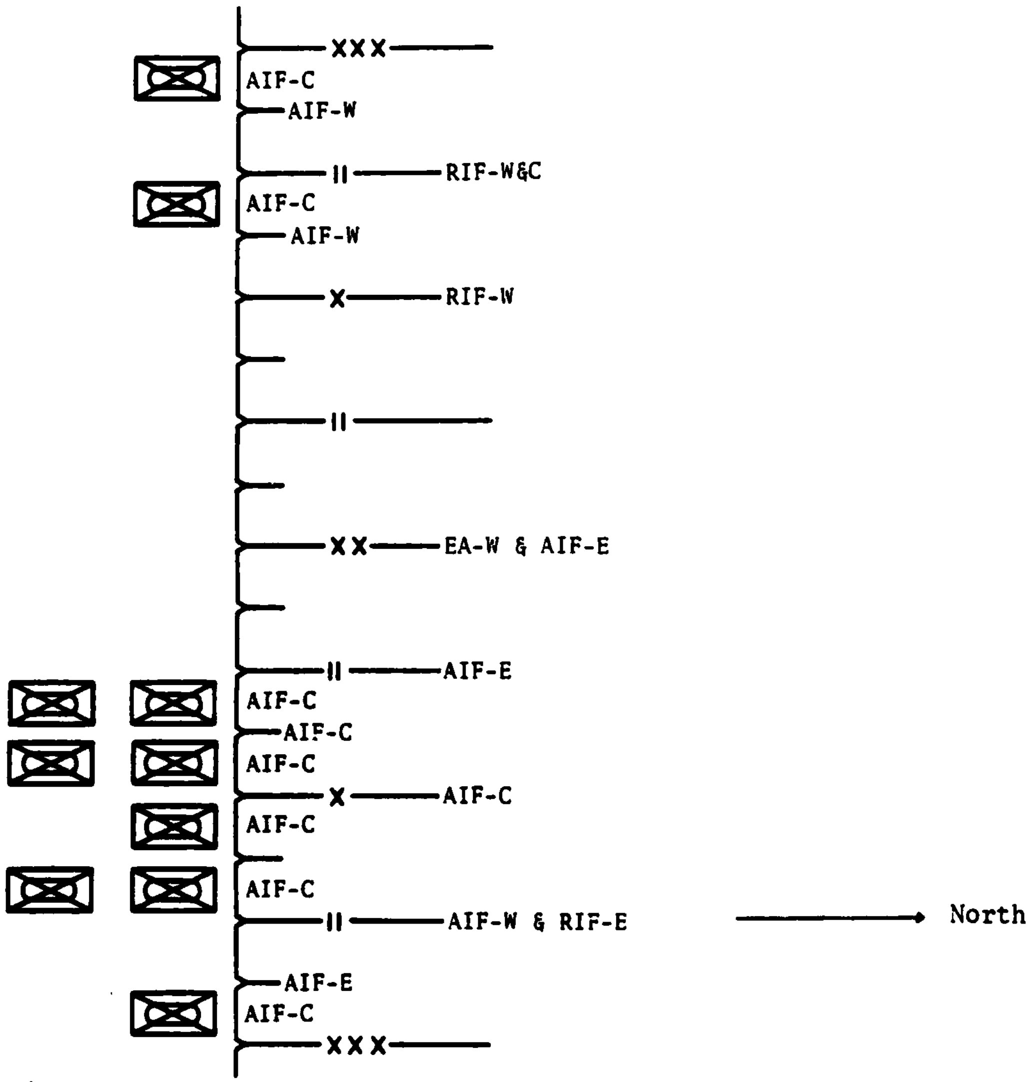

AIF-W - Attack In Force - West
AIF-C - Attack In Force - Center
AIF-E - Attack In Force - East
RIF-W - Reconnaissance In Force - West
RIF-C - Reconnaissance In Force - Center
RIF-E - Reconnaissance In Force - East

EA-W - Enemy Activity - West
EA-C - Enemy Activity - Center
EA-E - Enemy Activity - East

FIGURE 4
THE DECOMPOSITION OF INTELLIGENCE SUMMARIES BY SECTORS

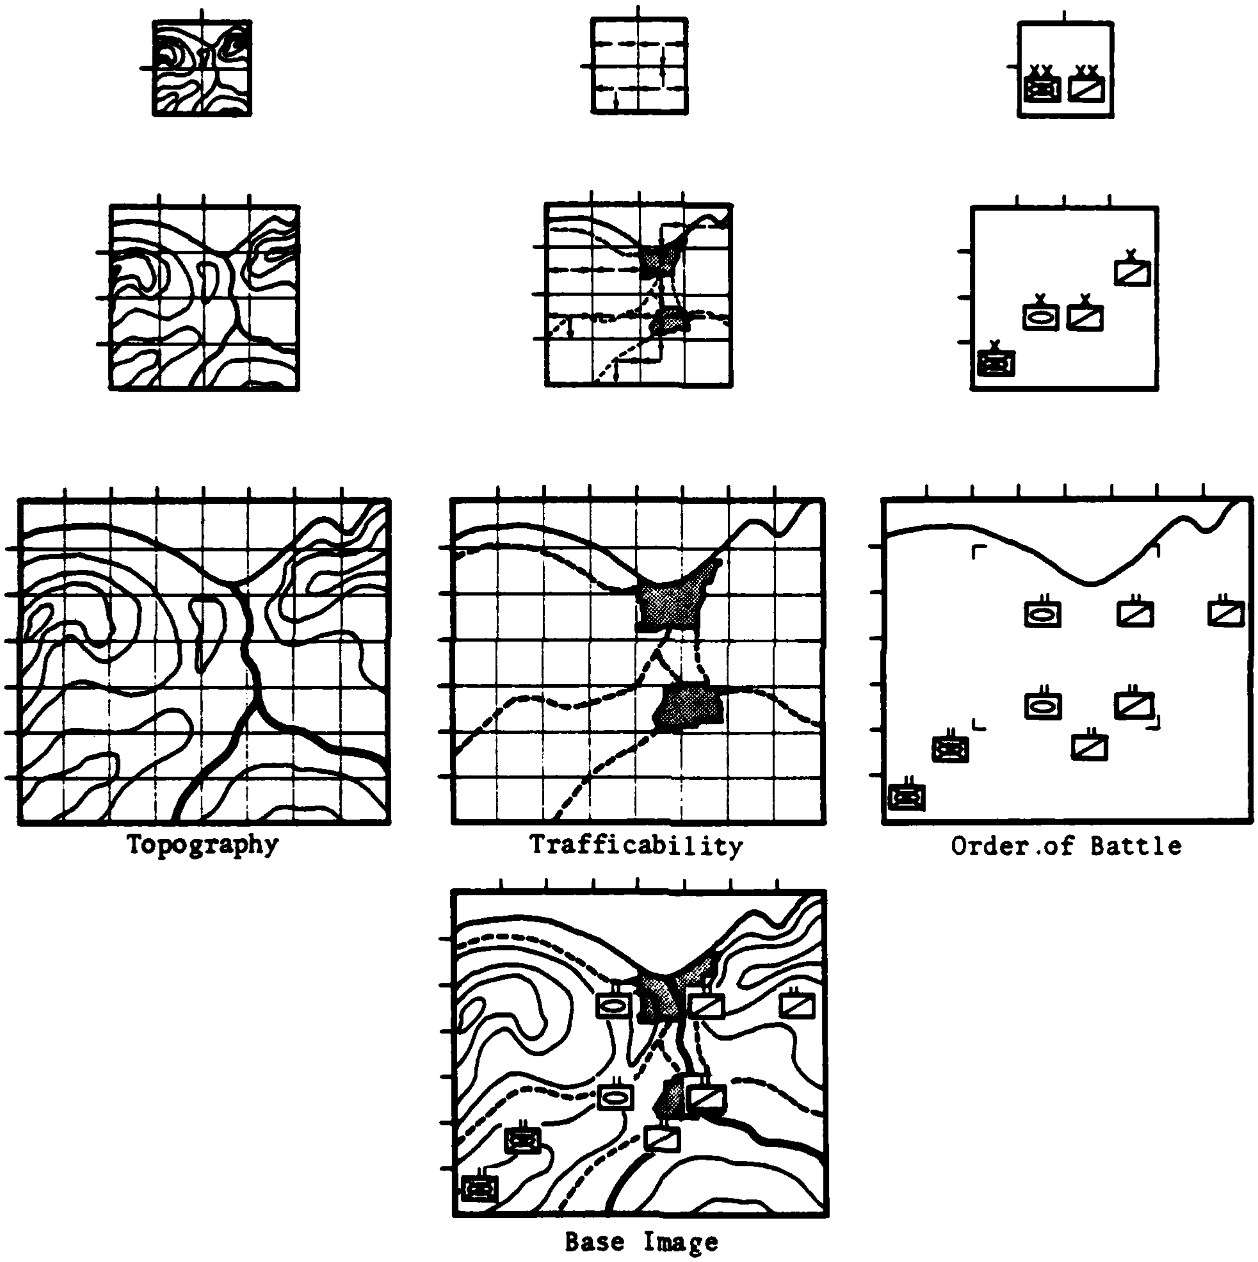

FIGURE 5
OVERLAYED DATA BASE REPRESENTATION

also facilitates the application of the more advanced AI techniques such as the meta-rules described above since the system can treat its knowledge directly as data. Frames also provide a structure for generalizing rules over instances and for efficiently implementing and using large bodies of rules. Large numbers of inference rules can be reflected in the default mechanism of the frame structure and reasoning at multiple levels of detail can be carried out in a straightforward manner. This considerably simplifies the rule hierarchy alluded to in the last section, allowing induction and generalization to be represented as just a stepping up in the hierarchy. These automatic inferencing and abstracting capabilities are especially attractive in light of the hierarchic nature of the threat that the fusion module must characterize.

3.2.2.2 Database Structure and Content. Conceptual clarity is served by segmenting the databases into logical categories of facts. The categories suited to our intelligence fusion approach are the perceived battle situation, terrain conditions, weather conditions, the disposition of friendly forces, the expected enemy order of battle, equipment characteristics, and system history.* We describe these briefly and in turn in the rest of this section.

The Sitmap

The perceived battlefield situation is represented in a system-generated and a system-maintained knowledge base of the current, derived knowledge called the situation map (Sitmap). The Sitmap is a frame with four slots, one for each quadrant of the battlefield. These slots either hold pointers to quadrant frames or are pieces of the battlemap (represented as picture element, or pixel, arrays), as illustrated in Figure 6. The arrays are marked with features and overlaid with report clusters that indicate enemy activity.

* Although all derived order of battle results are placed in these databases, they are an inadequate representation for higher level intelligence products such as, "the enemy will use tactical nuclear weapons." Such products will be stored with the rules that produced them.

```
Sitmap frame:
   quadrants     :   (S2 S3)
   reports       :   (MTI 30  ELINT 20  HUMINT 5)
   activities    :   (Radar 6  Artillery 10  Maneuver 12)

   S2 frame
      quadrants     :   (S23  S24)
      reports       :   (MTI 15  ELINT 10  HUMINT 5)
      activities    :   (Radar 6  Artillery 7  Maneuver 5)

      S23 frame
         quadrants     :   (S232)
         reports       :   (ELINT 7  HUMINT 5)
         activities    :   (Radar 5  Artillery 5)

      S24 frame        :
         quadrants     :   (S243)
         reports       :   (MTI 15  ELINT 3)
         activities    :   (Radar 1  Artillery 2  Maneuver 5)
```

S243 frame:

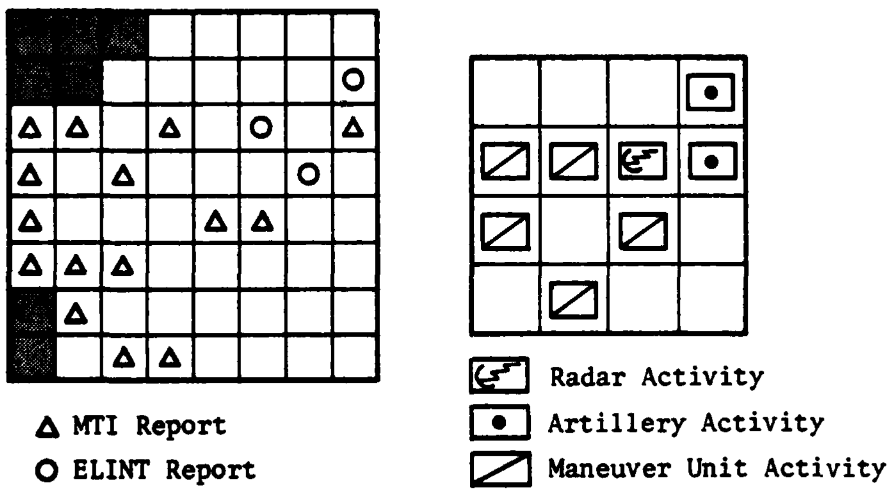

FIGURE 6
SITMAP REPRESENTATION

Terrain

Terrain strongly influences the disposition and deployment of battle forces. There is a strong connection between the expected patterns of enemy deployment from a doctrinal point of view and enabling or constraining characteristics of the environment. In order to develop a realistic representation of intelligence fusion, these environmental factors must be taken into account. This is done by explicitly stating the effect of terrain on war-making behavior in the terrain rules. The proper employment of terrain information is then facilitated by a hierarchical data structure registered the Sitmap. This is illustrated in Figure 7.

Weather

Weather will use the same organizational structure as terrain to make the environmental impact of weather on terrain easy to determine. The depth of the weather-frame will probably be much less than the terrain-frame, since weather events occur at a larger scale. This is illustrated in Figure 8, where Rain blocks that cover large corresponding areas of Figure 7 have varying effects on the forest, swamp, and agricultural cells of the terrain representation.

Friendly Dispositions

In order to 1) assess the potential threat of enemy actions on friendly activities, and 2) to determine what resources are available and appropriate to task for information needs, the system needs to have a representation of friendly force dispositions and collection assets. This database needs to be initialized by the other simulation modules (or by hand by the modeler) and kept up to date by the fusion module during the simulation.

The natural representation for the friendly forces is a hierarchial decomposition along lines of control (Figure 9). It will be convenient to supplement this with a representation of friendly units zones of responsibilities. The zones of responsibility will be further subdivided into areas recently covered by some sensor and those areas not recently covered by any sensor. This will make it possible to distinguish between a quiescent enemy and an absent one. The units collection capabilities will have to be represented to make appropriate tasking requests. For

```
Terrain frame:
   quadrants    :   (t1  t2  t3  t4)

   t1 frame:
      descriptor   :   agriculture

   t2 frame:
      quadrants    :   (t21  t22  t23  t24)

      t21 frame:
         descriptor   :   agriculture

      t22 frame:
         descriptor   :   forested area

      t23 frame:
         descriptor   :   forested area

      t24 frame:
         descriptor   :   agriculture

   t3 frame:
      quadrants    :   (t31  t32,  t33  t34)
```

t3 frame:
quadrants: (t31 t32 t33 t34)

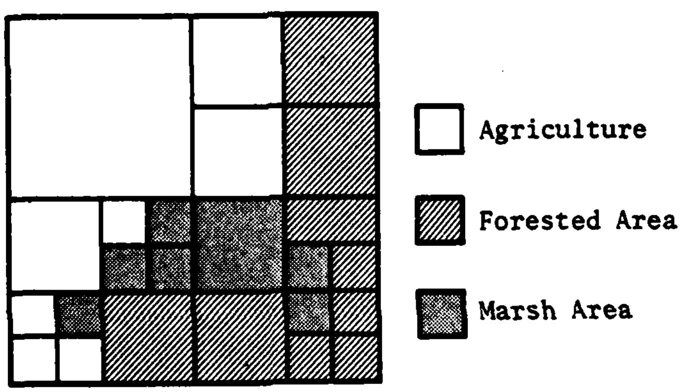

FIGURE 7
TERRAIN REPRESENTATION

```
Weather frame:
   default        :   clear
   quadrants      :   (W1  W3  W4)

   W1 frame:
      quadrants    :   (W13)

      W13 frame:
         descriptor   :   rain

   W3 frame:
      quadrants    :   (W33)

      W33 frame:
         descriptor   :   rain

   W4 frame:
      quadrants    :   (W41  W42  W43)

      W41 frame:
         descriptor   :   rain

      W42 frame:
         descriptor   :   rain

      W43 frame:
         descriptor   :   hail
```

R = Rain

H = Hail

Otherwise = Clear

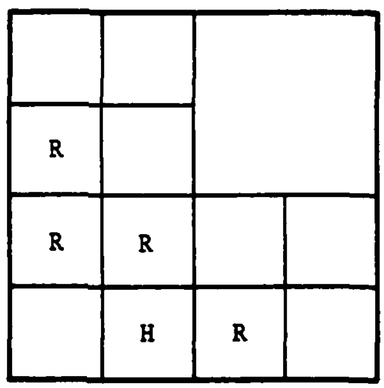

FIGURE 8
WEATHER REPRESENTATION

```
Friendly Dispositions Frame:
   quadrants   :   (friendly 4)

   Friendly 4 frame:
      unit            :   2nd-Division
      subordinates    :   (friendly 42  friendly 44)

      Friendly 42 frame:
         unit            :   2nd-Armored-Brigade
         superior        :   2nd-Division
         subordinates    :   (friendly 132  friendly 422)

         Friendly 132 frame:
            unit            :   106th-BN
            superior        :   2nd-Armored-Brigade
            subordinates    :   (friendly 1322  friendly 1322)

            Friendly 1322 frame:
               unit            :   B-Company
               superior        :   106th-BN
               enemy-facing    :
```

FIGURE 9
FRIENDLY DISPOSITIONS

instance, MTI will be requested in a sector only of MTI-equipped units who can cover the sector. The representation for friendly dispositions will follow the same region decomposition used in the Sitmap, terrain, and weather representations.

The friendly disposition database will be used to organize the fusion module's entire PIR/IR response task. Since PIR/IR are generated by specific units, the information relevant to that unit will be found in the unit's geographical locale. This is illustrated in Figure 4, where a characterization of enemy activity takes on very different forms depending on who is requesting it. Sometimes activities occurring far from the immediate scene of action have a bearing on the intelligence fusion task. Such is the case, for instance, with long range missiles. In order to represent this a "top.node" to which everything is relevant is required. The node is included in the disposition hierarchy to insure that enemy activity that is potentially relevant but far from any unit is taken into account.

The friendly dispositions represented in the fusion module are not to be understood as a full representation of the blue side of the model. On the contrary, blue is modeled in a minimal way to avoid redundancy. The skeleton presented is, as will be made apparent in the next section, just enough to support the fusion module's organizational structure. All other information about blue forces must be specifically requested from other modules.

Enemy Order of Battle

Part of intelligence analysis depends on expectations of enemy dispositions based on known enemy order of battle. We represent order of battle by a hierarchy of superior and subordinate unit frames keyed to the geographic representation of the terrain database and the friendly force sectors in which they appear. The representation closely parallels that of friendly dispositions, except that as many characteristics of the enemy unit as possible are attached to the unit frame.

The precise disposition of enemy forces is the nearly unachievable goal of a fusion center. In light of the importance of the intelligence mission and the enemy's attempts at deception, great care must be taken to preserve genuine information about the enemy and keep track of intelligence that is the result of guesswork. Therefore, among the characteristics attached to a unit's frame are its epistemic status, that is, whether it was observed, inferred from indicators, or deduced from other considerations. Often the existence of a unit in an area follows from EAC intelligence of its parent unit in the area. In this way information about enemy dispositions can be generated "top down," from parent unit locations (e.g. "expectations"), as well as "bottom up," from reports.

Equipment

A database of equipment types and characteristics is kept separately from equipment inventories of units to make general reasoning about technological capabilities possible. Classes of equipments contain sub-classes and instances of the class. As usual, equipments may inherit properties by membership in classes of equipment types.

System History

The fusion module's activities must be recorded because a) some of the findings of the fusion center may be overturned at a later time, and b) the fusion module must be able to learn from its mistakes. These two needs can be met by linked rules to the PIR. Antecedents that are established while trying to answer a PIR are "marked" with that PIR. When one of its antecedents is overturned, the effects on the PIR needs to be computed. If the resulting change in a PIR is sufficiently great, a message to that effect is sent to the module that originally requested the PIR, and the PIR record is marked as having been overturned. The marking of overturned PIR allows the system to focus on just those rule sets that have led to mistakes that is, overturned PIR/IRs. In this way, confidence can be dynamically redistributed through the system, and the affected rule set changed.

There are two other functions the System History may perform. 1) It will initiate the module's activity by having a set of standing PIR/IR's which the fusion module will immediately fulfill. This insures that critical developments will be reported to Force Control even if it did not explicitly request reports on such developments. 2) At intervals the System History can be reviewed and used to produce an Intelligence Summary. Such a capability is particularly useful in the training and war-gaming modes of operation.

3.2.3 The Inference Engine, Mechanics

The inference engine turns the static structures in the data and knowledge bases into result-generating processes. This is triggered by input from other modules in the IEW simulation model as illustrated in Figure 10. The inference engine is basically a pattern matcher that decomposes the incoming messages into strings that also occur in the fusion module's data or knowledge bases. The occurrences most often are partial matches. For instance, a request may read "How many enemy tank battalions face the 3rd Brigade." This may decompose to "(tank battalions/sector 2) + (tank battalions/sector 3)." The clauses in the transformed request may match elements in the database "tank battalions/sector 2=3" and "tank battalions/sector 3=5". The inference engine receives the non-matched, value-carrying part of the partially matched data base elements, adds them, and returns the message "There are 8 tank battalions facing the 3rd Brigade."*

Notice that the inference engine must translate an input from a simulation module into an internal process (e.g. "how many" = summing of counts) and also needs to translate a specification derived from another module's format into its own representation (e.g. "facing the 3rd Brigade", into the "sector" representation). The inference engine is in this respect data-driven. The input determines the procedure invoked. If the message is an intelligence report, for instance, an insertion or some other database updating operation is required.

The Inference Engine must know not only how to apply the operations that appear in the rules, but also how to determine the reliability of a statement that is the consequent of other statements of varying probabilities. It is adequate, in a rough approximation of such reasoning, to take the maximum of the probabilities of a group of OR-ed statements, the minimum of AND-ed statements, and the average of evidence statements. This follows the intuitions of the early fuzzy logic techniques.[20] However, advances in the Bayesian theory [21] and in the evidential analysis[22] offer better solutions. The best solution for the intelligence domain is

* The English sentences are illustrative only. No natural language interface to the fusion module . necessary in an automated environment. In a systemic model, a database management system (dbms) query translator will suffice.

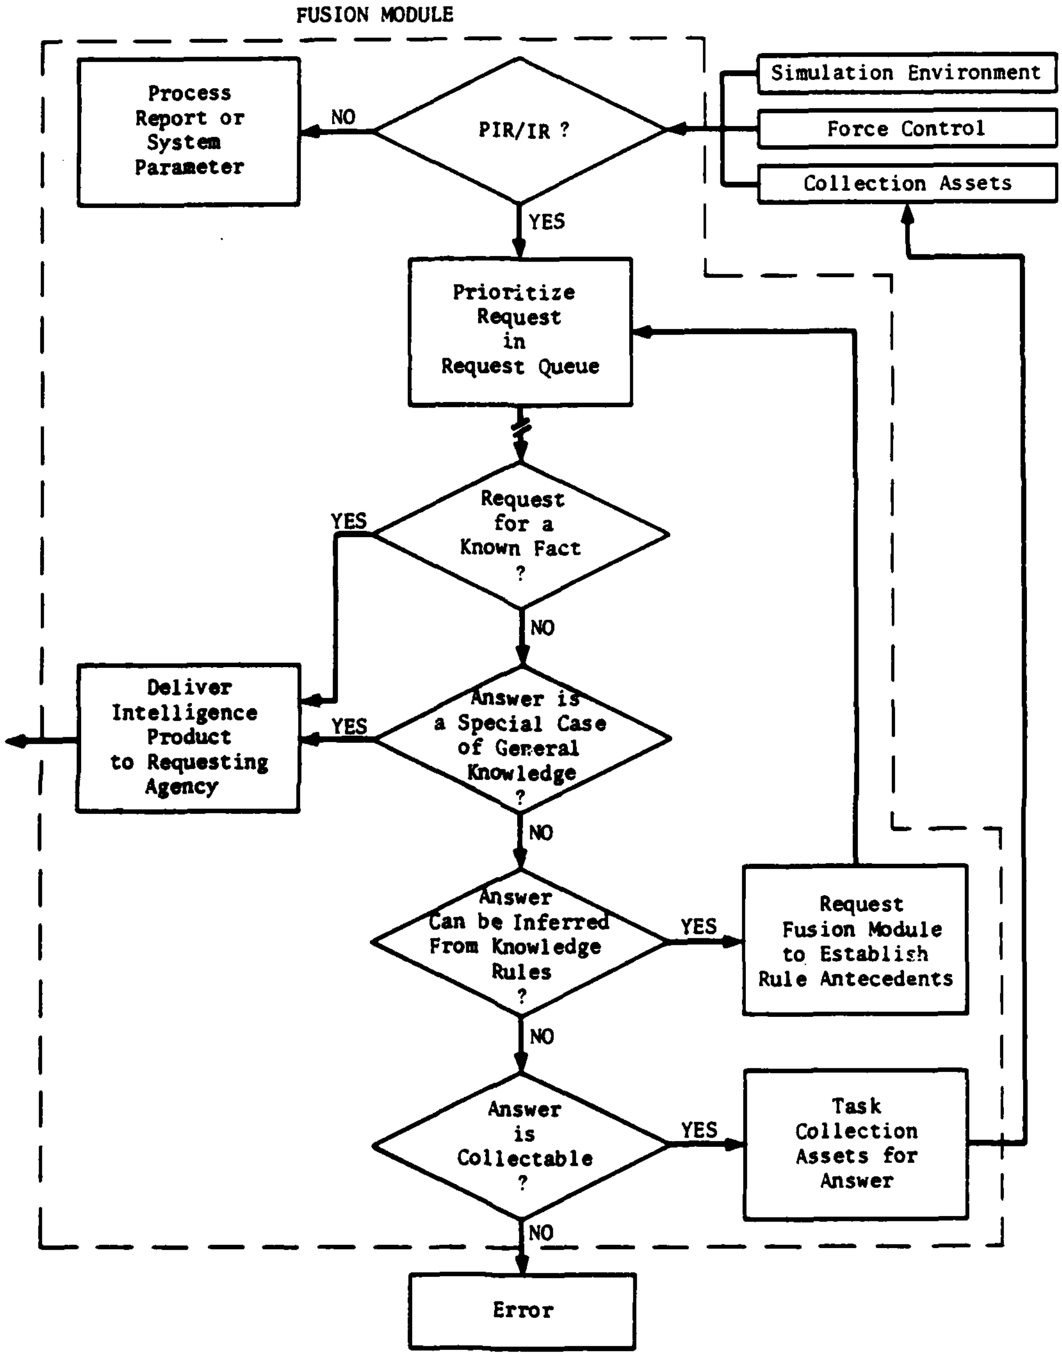

FIGURE 10
INFERENCE ENGINE FLOW OF CONTROL

explicit evidential reasoning.[23] This has the advantage of making justification transparent, but requires a greater investment in system development time. For the initial design the fuzzy logic approach is recommended. Implementation of this inference engine activity by meta-rules makes upgrades to the other techniques easier.*

The reliabilities of the statements upon which the intelligence product is grounded are not independent, but related according to the rules in which the statements appear. Suppose we are interested in whether the enemy is reinforcing and the following two rules appear in the knowledge base:

CI rule 4 (confidence - 75%)

<u>If</u> there is movement of additional troops toward the front
<u>or</u> increased traffic toward present positions
<u>or</u> new units have been identified in the combat zone
<u>or</u> additional CP's and supply and evacuation installations have been reported
<u>Then</u> the enemy is reinforcing the front

CI rule 5 (confidence - 80%):

<u>If</u> there is a new unit found in the battle area
<u>and</u> the unit is identified as (part of) a previously unseen unit
<u>Then</u> new units have been identified in the combat zone

* Different techniques require different additional pieces of information to be attached to the rules, so the combination strategy is not completely independent of the rule representation.

Suppose further that the reliabilities of the four antecedents of rule CI4 are .5, .5, .1, and .1, respectively,** and those of CI5 are 1.0 and .1. The probability that the enemy is reinforcing is determined by the maximum of the reliabilities of its supporting antecedents, because they are OR-ed, and is .5. If we now receive information that identifies a new found unit as part of a previously unseen regiment with high, say .9, reliability, we must recompute the reliability "new units have been reported in the combat zone." It's support is and-ed together, so the min (.9) of the reliabilities of its support is taken as the reliability of the statement. Now "the enemy is reinforcing the front" must be recomputed, since its support has changed. In fact, the max of its support is now .9, from the original new information about enemy deployments. In this way the information in propogated through the system to reflect changing beliefs about the battle situation.

The fusion module may not immediately "know" the answer to the request. This will be the case when the request is a "higher level" information that only appears in the rule consequents and (possibly) antecedents. The inference engine then turns to the knowledge base appropriate to the request and matches the transformed request statement to the conclusions or "consequents" of the rules in the knowledge base.

More than one rule may match a request. When this occurs the matching rules are said to conflict. In this case, the Inference Engine needs to be able to reason about the rules in the conflict set to prioritize them. This can be done implicitly with an intrinsic ordering on the entire rule set or by meta-rules which provide ways of explicitly reasoning about conflict resolution. Explicit representation in meta-rules (as in meta-rule 3 in Section 3.2.1.5) is recommended.

The above example is a best-case situation in that a data base had values for the matched facts. This will not always be true. When the required information is missing, the inference engine must "be informed," and must respond accordingly. In the fusion module this means the inference engine will use "default reasoning" to

** The reliability of "new units have been identified in the combat zone" is derived from the min of 1.0 and .1. The other statement's reliabilities would similarly be derived from the reliability of the evidence supporting them.

make a good guess at what the values might be or it will task the IEW collection module to provide the missing information. It may also invoke knowledge about the reliability of its default reasoning and the probable cost in time and attrition of the tasking effort in order to decide which action to take.

We will illustrate these ideas with a small example. Suppose the request "Can the enemy 103rd Regiment make a deep penetration?" is received and it fails to match in the database. The inference engine turns to the knowledge base and, lets say, finds a rule that states:

CI rule 6 (confidence - 70%)

<u>If</u> the unit is well supplied
<u>and</u> the unit has air support
<u>and</u> the unit has a decisive advantage in armor
<u>Then</u> the unit can make a deep penetration.
(with confidence 70)

Now the inference engine tries to establish the antecedent premises by matching their forms to the database or knowledge base contents exactly as it did with the original intelligence request. It literally asks itself if it knows or can determine each antecedent in turn. When it tries to determine "the unit has a decisive advantage in armor," it might find another rule that states:

CI rule 7: (confidence - 85%)

<u>If</u> (3x the number of armored units of size x in the unit)
is greater than
(2x the number of armored units of size x facing the unit)
<u>Then</u> the unit has a decisive advantage in armor

The inference engine is back to the level of the facts in its database and can proceed to determine the values of those facts or the tasking required, as described above.

3.2.4 <u>System Parameters</u>

The previous sections describe a methodology that can be employed to accomplish intelligence analysis as well as model it. In the modeling environment, however, a number of additional factors must be considered in order to faithfully represent human analysts. These are the considerations that derive from the impact of the battlefield environment on a fusion center. This section will indicate how that impact can be reflected in the rule-based fusion module.

There are four main effects that must be modeled and two main ways of modeling them. The effects are weather, attrition due to combat, degradation due to transit, and fatigue due to continuous operation. The two approaches to implementing these effects are to change the input before the rule-base processes it or to change the rule-base that is used to process the input. The paragraphs below describe the implementation of each effect.

We have seen that some aspects of weather are best modeled explicitly in a rule-base. Other aspects are more conveniently handled through system parameters. Sensors are often sensitive to atmospheric effects and can be adjusted directly by changing the error ellipses and reliability ratings of the sensor reports. Increased noise in communications traffic can be simulated by randomly dropping some messages or by corrupting the message format. Since the fusion module operates using a formal input grammar, any noise that is added to the input will cause it to be incomprehensible to the fusion module.

Down-time of automated data processing equipment and casualties to skilled personnel are best modeled using a combination of approaches. Since these elements play very specific roles in the fusion center, the rules which involve or reflect those roles must have their reliabilities altered or, perhaps be dropped altogether. However, computers also fulfill the general roles of data storage and high-volume processing. The loss of the computer facility therefore impacts the fusion center's ability to handle volume effectively. This is most easily modeled by setting a required reliability threshold of reports that are processed. When this is done, the number of input reports that are used by the fusion module decreases and more default reasoning is invoked. The quality of the intelligence product can be expected to drop accordingly.

Displacement of a fusion center to another location results in degraded performance for roughly the same reason as a computer failure and personnel loss, so the same threshold technique used to model combat casualties in men and equipment can be used for modeling effects due to transit. The only difference is that the effect will be of shorter duration.

Fatigue due to continuous operation will affect the behavior of a fusion center.* Under stress and fatigue, subtle points are missed and intelligence suffers. In the idiom of rule-based systems, the rules of thumb used by the combat intelligence staff change in the direction of coarser reasoning. This suggests changing the rule-base itself through the application of meta-rules, or even applying a special rule set for approximating these effects. However, this can again be approximated by raising of the reliability threshold of input reports. Thus only the most obvious inferences will remain visible to the fusion module.

* This is mitigated by the degree of automation at the fusion center.

4.0 EXAMPLES

In this section we develop an example to illustrate the methodology. We simulate a series of events consisting of the message inputs sent to the fusion module from other entities in the model environment, the outputs from the fusion module, and the system messages that come from the scenario operating environment. These events indicate the kind of activity this methodology is capable of modeling.

The example is derived from FM30-5[24], the field manual for Combat Intelligence. It is a walk-through of Essential Elements of Information (1) and (2) and Other Intelligence Requirements (1) in "Example of a division intelligence annex," FM30-5, Appendix N, Section 2. The critical indicators used in the example are taken directly from "Typical Intelligence Indicators," FM30-5, Appendix T. It must be emphasized the we depend heavily on our natural language understanding to draw the inferences in the example. In the actual automated model, the PIR/IR language has to be explicitly factored into knowledge items that ultimately translate into the "observable" data elements of the database.

The example is presented as an annotated scenario of 16 events beginning at 0900 on September 10 and ending at 1500 on the same day, and occurring in a contested area of strictly hypothetical existence. Events have the following format: The number of the event in the event sequence is given along with its type and descriptor (e.g. Input: SYSMSG 002 is a system Message proper in bullit form.) The last section of the event is the comment section, which annotates what has transpired.

Event 1

Input: PIR 001 (10 0900 Sept CORPSHQ)

- •• Will aggressor reinforce his forces along the FLOOD River before the time of attack (101800 SEPT)?
- • If so, when, where, and with what forces?
- • Special attention to the mechanized regiment and the medium tank regiment in the vicinity of BURG.

Comment:

At 0900 on September 10 an intelligence request and two sub-requests are receive from Corps Headquarters. The primary mission is to determine if the enemy is being reinforced in a particular area. The system understands reinforcement as follows (derived from FM30-5, Appendix T, Section T-6).

REINFORCEMENT:

- Movement of additional troops toward the front
- Increased traffic toward present position
- Identification of new units in combat zone
- Additional command posts and supply and evacuation installations

The reinforcement is further broken down as:

- Movement of additional troops toward the front
  - MTI in direction of front
  - COMINT
  - HUMINT
  - EAC Information
  - cluster, pattern, unit

- Increased traffic toward present position
  - MTI
  - COMINT
  - HUMINT

- Identification for new units in combat zone
  - pattern, unit
  - EAC Information
  - HUMINT and CLUSTER

- Additional CPs and supply and evacuation installations

We assume that the only information the fusion module has regarding enemy reinforcement is an EAC indication to that effect. We also assume that that is not enough to state categorically that reinforcement is occurring. It has no other data bearing on the request.

Event 2
Output: PIR 001 collection tasking

MTI report clusters toward FLOOD River and BURG?
(100930 SEPT Fusion)

Comment:

The fusion module has been unable to derive an answer to the PIR from its present knowledge and data. It tasks the collection module's MTI assets to provide the most important mission information. Implicit in the fusion module response is an understanding of what is adequate to establish reinforcement (e.g. three out of four antecedents established). For simplicity we assume an MTI report would be sufficient.

Event 3

Output (to self)/Input:

PIR 001 set demon = (PIR 001 101300 SEPT 101800 SEPT Aggressor is reinforcing lightly along the FLOOD; Receiving sub-regimental troop level N of FLOOD; No information about reinforcements in BURG vicinity)

Comment:

The fusion module sets a demon (by sending itself a special message). The demon holds the PIR as it was determined before tasking. It is to wake at 1300, September 10, if PIR 001 has still not been determined by then.

The fusion center needs to be kept apprised of events in the modeling environment so that the demon facility will work. This can be efficiently accomplished by putting a "demon wake-up demon" in the event queue of the global model, whenever a demon is posted to the fusion module. Then, when the demon wake-up demon becomes the scheduled event in the model, it sends a dummy message to the fusion module. The fusion module then prioritizes the input, notices its demon message ("wakes" it), and takes the action the demon indicates. This way the demon will fire even if no PIR/IR or reports are sent to the fusion module before the demon's wake-up time.

Event 4
Input: PIR 002 (101100 SEPT Corps HQ)

- •• Will Aggressor employ tactical nuclear weapons against us?
- • If so, when, where, how many, of what yields, and by what delivery means

Comments:

Another PIR is received. The fulfillment of the request depends on indicators of use of nuclear weapons, found in FM30-5, Appendix T, Section T-8.

Nuclear Weapons Use

- Location of missile and/or free rocket units within striking range
- Use of missiles and/or free rockets with high explosives
- Location of very heavy artillery within supporting distance of front lines
- Registration with very heavy artillery
- Special or unusual activity by frontline troops
- Limited withdrawal of frontline troops without apparent tactical reason
- Sudden and energetic digging in enemy areas
- Large concentrations of radios, radar, and other electronic equipment located in vicinity of suitable sites for guided missile launching

Event 5
Output: (PIR 002 Corps HQ

Aggressor will employ tactical nuclear weapons with low probability

Not before 130000 SEPT
Probably between Hill 438 and Hill 439
One to three missiles of
One megaton each
Free rockets based between BERGEN Ridges

(101200 SEPT Fusion))

Comment:

In this case the fusion module has inferred the PIR from previous information.

Event 6
Output: (*DEMON* (PIR 001 Corps HQ)
Aggressor is reinforcing lightly along FLOOD

Receiving sub-regimental troop level North of FLOOD

(101300 SEPT Fusion))

Comment:

Since no MTI reports have been received by 1300, the demon set in event 3 fires, sending the message stored with it to Corps HQ.

Event 7
Input: SYSMSG 001 (101300 SEPT SYSTEM)

- Weather change - thunderstorms south of BURG

Comment:

A system message is received. The system stores information about the environment, in this case that there are thunderstorms in the area of the fusion center.

Event 8
Input: SYSMSG 002 (101315 SEPT SYSTEM)

- Fusion Center Computer Facility Down until 101400 SEPT

Comment:

A second system message (SYSMSG) is received. The thunderstorm has caused the computer facility to fail. The effect is modeled by a change in the acceptable threshold for message reliability (Section 3.4).

Event 9
Output (to self)/Input

SYSMSG 002 (101315 SEPT FUSION)

- Set Previous to Report Processing Reliability Threshold (RPRT)
- Set RPRT = (1 - (.15 * Previous))
- Set DEMON = (sysmsg 002 101400 SEPT Reset report processing reliability threshold to Previous)

Comment:

The effect of the system message is implemented by the same mechanism that was used for demon planting. In fact, it sets another demon to restore the computer operations at 1400.

Event 10

Input: IR 003 (101330 Corps HQ)

- Will Aggressor continue to defend in his present position?
- If so, how will he organize his forces on the ground, and with what troops?
- Special attention to locations and activities of reserves and vulnerability to nuclear attacks.

Comment:

At 13:30 an IR concerning enemy defensive posture is received. Indicators of defense (FM30-5, Appendix T, Section T-4) are:

DEFENSE

- Preparation of battalion and company defense areas
- Extensive preparation of field fortifications
- Formation of anti-tank strong points
- Attachment of additional anti-tank units to frontline defensive position
- Preparation of alternate artillery positions
- Employment of moving artillery
- Large tank units located in assembly areas to the rear
- Preparation and occupation of successive defense lines
- Presence of demolitions, contaminated areas, obstacles and mine fields
- Deployment of mechanized units on good defensive positions
- Dumping ammunition and engineer supplies and equiptment and fortifying buildings
- Entrenching and erecting bands of wire

Event 11
Input: Report 001 - (101345 HUMINT)

- Dispersal of tanks and SP guns to forward units on FLOOD
Reliability = .75

Comment:

Report disregarded because below present processing threshold. Notice that it would have indicated the enemy is preparing offensive operations.

Event 12
Output: (IR 003 - Corps HQ

Enemy intends to defend present positions
mech inf regiment and medium armor in reserve near BURG
mech inf regiment HILL 581 - vulnerable (concentrated)
mech inf regiment on FLOOD

(101400 SEPT FUSION))

Comment:

IR003, event 10, is answered.

Event 13
Input: Report 002 - (101415 HUMINT)

- Enemy conducting drills and rehearsals near BURG Reliability = .70

Comment:

Report increases the reliability of PIR001 (reinforcing) and decreases that of IR003, but not enough to overturn results. But it does change IR003 enough to generate a collection tasking for confirmation. This is because the report is an indicator of an attack posture, rather than the defense posture reported in IR003.

Event 14
Output: (IR003 collection tasking

- Artillery positions well forward and concentrated anywhere along FLOOD?
  (101430 SEPT FUSION))

Comment:

The indicators for an attack are, among other things, rehearsals and drills in rear areas and artillery positions well forward and concentrated. The fusion module is now trying to establish that the enemy will attack because it just received an indicator to that effect, and because it previously reported that the enemy was in defensive posture. It has determined the most important indicator to be the concentration of artillery forward, and is tasking to establish whether or not this is the case.

Event 15
Input: Report 003 - (101445 COMINT)

- Artillery positions well forward and concentrated along FLOOD North of BURG; new tank unit near HILL 581

Comment:

Report changes PIR 001 and IR 003 responses, since enemy is probably planning an attack and (therefore) bringing more forces to bear along the FLOOD North of BURG.

Event 16
Output: (PIR 001, IR 003 - Corps HQ)

- •• Enemy is reinforcing along FLOOD Northe of BURG probably before attack.
  - • with mechanized regiment and medium tank regiment and artillery
- •• Enemy will not continue in defensive role North of BURG
  - • will attack across FLOOD with two mechanized infantry regiment, medium tank regiment+, and artillery support.
  - • mechanized infantry regiment on HILL 581 will cease to be vulnerable to NBC due to dispersion by 101600.

Comment:

Since previous results are overturned, the fusion module looks in its PIR history and traces the affected requests. It finds that PIR001 and IR003 (but not PIR002) are affected and sends updated information to (the union of) their sources.

5.0 ISSUES

There are a number of issues relating to implementation and system development of knowledge-based systems that are germane here. These are known collectively as Knowledge Engineering (KE). We will briefly discuss the KE areas of knowledge acquisition, system completeness, documentation and consensus, efficiency, and host system requirements in the sections below.

5.1 Knowledge Acquisition

Knowledge-based systems depend for their success on a large data base of domain-specific knowledge. This knowledge base is built up incrementally and iteratively. Since the rule-based formalism is modular, knowledge may be added to the system, tested, and incorporated, modified, or deleted.

Frequently a non-specialist in the field of interest is responsible for building the knowledge base. A skeleton of the knowledge base can often be developed by common sense and the judicious application of ideas from textbooks on the subject. If the person responsible for imbuing the system with this knowledge is also a programmer familiar with the system implementation, extensive software support may not seem necessary. However, to move the system past the demonstration stage to a successful and mature knowledge-based prototype requires interaction with the real experts of the problem domain and requires software support for long-term involvement of non-programmers in knowledge editing and debugging.[13,25] This software development must be factored in to any plan to implement a fusion system such as described above.

5.2 System Completeness

A different problem in Knowledge Engineering is guaranteeing the completeness of the rule-base. A complete rule-base is one whose rules cover the situations that can be expected to arise in the system's domain. The best way to approach this is to attempt a top-down decomposition of the domain of interest. The intent is to make sure that there is at least one conceptual entity for each branch of the decomposition. Assuming that all relevant entities have been specified in a taxonomy, a rule checker can be constructed that tests for completeness among the concepts. It determines, for each entity and by strictly syntactic methods, whether it appears in the rule set, is derivable, or appears in contradictory rules. It can also keep track of conjunctions in the rule antecedents. The system can then inquire whether a new rule is correct if it contains fragments of conjunctions it has seen.[26]

A completely different approach is required if one does not with confidence believe that all relevant concepts appear in the system. In this case learning techniques that can generate new concepts are required. The inductive method[14] will not work under these circumstances since it just generalizes on existing entities. However, learning algorithms exist (e.g., the genetic algorithm[15]) that were designed precisely to automatically create new entities. The latter is the right approach for evolving the system over the long run.

5.3 Documentation and Consensus

When a knowledge base is developed, an audit trail of the growth of the rule set is needed. Documenting the person who added a rule, the time it was added, and the reasoning behind adding it makes it easier to gather confidence for the rule-base and makes changes to the rules more focused.

A related issue is gathering a consensus for a rule-base. This is very important for system acceptance, but some extra effort will probably be needed to achieve it. It may be necessary to run a set of experts through a set of example scenarios. Any rule that is used without complaint or comment by an expert gains that expert's tacit approval. Rules are then annotated with the results of the session or revised as a result of the comments. A distinction must be maintained between a rule's popularity and its success. The rule's efficacy in achieving desired results will always be the first consideration in its evaluation.

A rule may achieve the right results for the wrong reasons. Since we are interested, in this application, only in the operational consequences of the rules, the revision must be focused on the rules that lead to the firing of the rule in question. If the whole set of context rules also achieve the right results "for the wrong reasons," we do not revise any of them. Any revisions that are made are expected to suggest further corrective revisions until the rules achieve the right results for the right reasons.

5.4 Efficiency

Rule-based systems tend to become slower the more rules they contain, whereas humans tend to be able to perform tasks faster as they develop expertise in the task. This indicates that much work remains to be done before we really understand how expertise should be acquired and applied. However, improvements have been made in expert system efficiency to greatly reduce turnaround time for knowledge-based system outputs. Among the improvements are the application of better search strategies to the rule-base[26], new internal representations for the rules'[27] rule-base compilation, and decomposition of the solution space of the problem into subproblems

5.5 Host System Requirements

The system described is intended for a mini-computer. Although there is precedent for fitting a capable production system on a personal computer,[28] many of the features that make this architecture attractive for analysis may have to be jettisoned due to lack of space.

There is also precedent for implementing systems similar to this fusion system in languages such as PASCAL[14,28] and FORTRAN.[29] These languages are extremely cumbersome for many of the techniques that are appropriate for this task. However, the LISP family of languages provide the required symbolic, dynamic, and flexible constructs for implementing the above methodology, and we recommend it for this module. Moreover, small and fast workstation LISP computers are being marketed that will make the choice of LISP practical even for application requiring speed and portability.

6.0 SUMMARY

We have presented an AI-based methodology for modeling intelligence fusion. The methodology depends on recognizing two different kinds of demands on the fusion module - integrating intelligence reports and responding to intelligence requests. We employ forward chaining through a rule-base for the former, and goal-directed backchaining for the latter. Both facts and rules are represented in frame structures, making it possible to apply rules to rewrite rules and to explicitly manipulate the control strategy. In this section we review our methodology in light of the problem definition of Section 2.

The Functional Area Representation Objectives state that fusion is to use "sensor reports of all types along with terrain and weather data to determine enemy location, strength, and intent".[5] SIGINT, MTI, IMINT, and HUMINT sensor reports of all kinds are used in the forward-chaining portion of the methodology to determine enemy strength and location. Enemy intent is developed by back-chaining through the Enemy Characteristics rule base to enemy unit locations and strengths. Terrain and weather enter into intelligence considerations through the integral construction of the Sitmap, terrain, and weather frames. Weather was also discussed as a system parameter. The IEW FARO's enumerate the effects of combat, the environment, movement, and other functional areas on the fusion center. Preprocessing and thresholding techniques are used to reflect these effects in fusion center behavior.

The Model Requirements Definition Document[6] describes the fusion module behavior in terms of an Input-Process-Output template covering collection, single-source correlation, and multi-source analysis. Collection is partially covered in the fusion module because the methodology decomposes input into recognizable collection requests in its goal-directed back-chaining part. The collection task proper remains outside of our system, however. Single-source correlation is achieved in the cluster analysis in the forward-chaining part of the method. The Multi-Source Fusion centers as described in the MRDD are actually represented by both the pattern rules segment of the forward-chaining component and the back-chaining PIR/IR processes in our system. The dissemination of intelligence products is not a responsibility of the fusion center, therefore, intelligence is directed only at the requesting agencies.*

The problem is characterized, from the computational point of view, by incomplete information, unreliable information, time-varying data, and by implicit structures in the data. The hierarchical frame representation of objects and entities of interest to the intelligence endeavor provides a natural solution to the incomplete data and implicit structure problems. Reliability is explicitly estimated, dynamically adjusted, and accounted for in the continual updating of data items in rules. An evidential reasoning approach is the recommended solution for the reliability assessment problem. The dynamic aspect of input data is modelled by the dynamic adjustment of reliabilities and the dependency structures of system knowledge that is an integral part of rule-based systems.

The constraints on the fusion module from the modeling point of view have their greatest effect on the fusion methodology implementation. Of the time compression, fidelity, and justifiability issues relevant to modeling, only the last pointed directly to a rule-based system. The speed-up of rule-based system execution time can be achieved in a number of ways, some of the most interesting of which employ learning algorithms to generalize and streamline the rule base. The fidelity of a rule-based model requires iterative building, testing, and modification until the system converges to the required specifications. This requires a software support environment in which rule editing and debugging is a planned endeavor.

* If this is found to be too limiting, an output buffer could be monitored by a "dissemination expert" and multiple copies of a product dispatched to the relevant agencies.

The fusion methodology we have presented is an application of well-tested software technology to the difficult problem of intelligence fusion. With direct implementation and proper knowledge engineering this technique will result in an improved model of human intelligence activities on the battlefield. With the features of spatial and temporal data structures, evidential reasoning, and the infusion of intelligence expertise, this technology can result in an extremely versatile simulator/trainer/wargamer of high capability and utility to the Army modeling community.

APPENDIX

Sources for Knowledge Based and Data Bases

In this appendix we note sources for the knowledge bases and the data bases on which the fusion module depends. It must be emphasized from the outset that this provides only the coarsest of guidelines for developing the rule and databases for the fusion center module. There are two reasons not to depend wholly on the sources described here. The first and most important is that there are broad differences between the procedures found in the FMs and practice on the field of action. The second is that the author is considerably more expert in AI technologies than military intelligence operations and the diverse sources upon which it draws. This list should by no means be construed as exhaustive.

Knowledge bases:

Report Processing Rules:	The rules for processing intelligence have been evolved interactively in the development of ANALYST. These may be obtained from R. P. Bonasso, the MITRE Corporation.
Terrain rules:	The rules for terrain processing may be extracted from FM30-5, Military Geographic Intelligence (Terrain),[31] and FM30-1CA, Special Applications of Terrain Intelligence.[32]
Weather rules:	Weather report processing rules may be developed from FM31-3/AFM105-4, Weather Support for Field Army Tactical Operations.[33]
Critical Indicators:	The critical indicators of enemy behavior were derived from FM30-5, Combat Intelligence,[24] and FM30-102, Opposing Forces Europe.[34] The threat we have considered is limited to Soviet and Warsaw Pact forces.

Meta-rules: The meta-rules appropriate to the fusion problem are indicated in Section 3.2.1.5. They will be evolved as the system is implemented. Guidance appears in the references found in the meta-rules section.

Data Bases:

Sitmap: The Sitmap is system built and maintained

Terrain: The aspects of terrain required for intelligence fusion are found in 66000-ASupR, Intelligence Preparation of the Battlefield.[35] The terrain data for actual battle areas is obtained from the Defense Mapping Agency.

Weather: The current weather picture is dynamically maintained by the system according to reports and weather rules.

Friendly Dispositions: The dispositions of friendly forces is initialized and maintained by the modeling system environment.

Enemy Order of Battle: The basic structure of enemy forces is delineation in FM30-103, Aggressor Order of Battle.[36] The local order of battle is partially initialized (previous intelligence) and evolved during the simulation.

Equipment: Equipment characteristics are contained in FM30-102, Opposing Forces Europe.

System History: This database is created and updated during a simulation run.

GLOSSARY

ADA	Air Defense Artillery
AI	Artificial Intelligence
AMIP	Army Model Improvement Program
AMMO	AMIP Management Office
ARI	Army Research Institute
ASPS	All Source Production Section
BDE	Brigade
BN	Battalion
C^2	Command and Control
CMDS	Collection Management and Dissemination Section
COMINT	Communications Intelligence
COP	Command Observation Post
CORDIVEM	Corps and Division Evaluation Model
CP	Command Post
EAC	Echelon Above Corps
ECM	Electronic Countermeasures
EEI/OIR	Essential Elements of Information/Other Information Requests (obsolete)
ELINT	Electronics Intelligence
EMP	Electro-Magnetic Pulse
EWS	Electronic Warfare Section
FAS	Field Artillery Section
FAROs	Functional Area Representation Objectives
FAS	Field Artillery Section
FLOT	Forward Line of Troops
FM	Field Manual
G2	Designation of division or corps intelligence staff
G3	Designation of division or corps operations staff
HUMINT	Human Intelligence
HVT	High Value Target
IED	Imitative Electronic Deception
IEW	Intelligence and Electronic Warfare
IMINT	Imagery Intelligence
INSCOM	Intelligence and Security Command
IPB	Intelligence Processing of the Battlefield
KE	Knowledge Engineering

GLOSSARY

(CONCLUDED)

LRRP	Long Range Reconnaissance Patrol
MGT	Management
MOPP	Mission Oriented Protective Posture
MRDD	Model Requirements Definition Document
MSN	Mission
MTI	Moving Target Indicator
NBC	Nuclear, Biological, Chemical
OB	Order of Battle
PIR/IR	Primary Information Requirements/ Information Requirements (formerly EEI/OIR)
REMS	Remotely Monitored Sensors
R/S	Reconnaissance and Surveillance
S2	Brigade Intelligence Officer
SIGINT	Signals Intelligence
SNR	Signal to Noise Ratio
SYSMSG	System Message
TCAE	Technical Control and Analysis Element
TRADOC	Training and Doctrine Command
TREE	Transient Radiation Effects on Electronics

REFERENCES

1. U.S. Army Research Institute and U.S. Army Intelligence and Security Command, "Investigation of methods and techniques for intelligence analysis - a joint ARI/INSCOM project", briefing materials.

2. Lockheed Missiles and Space Company, Inc., The Determination of Essential Elements of Information for an Army Corps, LMSC-D664802, November 1978.

3. Special Study Group, Department of the Army, Review of Army Analysis, April 1979.

4. Xerox Palo Alto Research Centers, The Organization of Expert Systems: A Prescriptive Tutorial, VLSI-82-1, M. Stefik et. al., January 1982.

5. The MITRE Corporation, Functional Area Representation Objectives for the Corps/Division Evaluation Model, MTR-82W00175, R. P. Bonasso, J. R. Davidson, P. K. Groveston, R. O. Nugent, October 1982.

6. The MITRE Corporation, IEW Model Requirements Definition Document, MTR-83W00130, J. R. Davidson, P. K. Groveston, June 3, 1982.

7. RAND, An Investigation of Tools for Building Expert Systems,R-2818-NSF, D. A. Waterman, Frederick Hayes-Roth, June 1982.

8. National Bureau of Standards, An Overview of Expert Systems, NBSIR-82-25, W. B Gevarter, May 1982.

9. The MITRE Corporation, ANALYST: An Expert System for Processing Sensor Returns, MTP-83W00002, R. P. Bonasso, 1983.

10. Lenat, D. B., "The Use of Heuristics to Guide Learning by Discovery," September 1981.

11. Michalski, R. S., "Pattern Recognition as Rule-Guided Inductive Inference", IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. PAM 1-2, No. 4, July 1980.

12. Aikins, T. S., "Prototypical Knowledge for Expert Systems", Artificial Intelligence, Vol. 20, No. 2, February 1983.

13. Van Melle, W., "A Domain-Independent System that Aids in Constructing Knowledge-Based Consultation Programs," Ph.D. Thesis, Stanford University, June 1980.

REFERENCES

(CONTINUED)

14. Michalski, R. S., R. L. Chilansky, "Learning by Being Told and Learning from Examples: An Experimental Comparison of the Two Methods of Knowledge Acquisition in the Context of Developing an Expert System for Soybean Disease Diagnosis," International Journal of Policy Analysis and Information Systems, Vol. 4, No. 2, 1980.

15. Holland, J. H., Adaptation in Natural and Artificial Systems, University of Michigan Press, 1975.

16. Slagle, J. E. Halpern, "An Intelligent Control Strategy for Computer Consultation", Naval Research Lab Memorandum Report 4789, April 8, 1982.

17. Minsky, M., "A framework for representing knowledge", Massachusetts Institute of Technology Artificial Intelligence Laboratory, AI Memo 306, 1974.

18. Antonisse, H. J., "Image Segmentation in Pyramids", Computer Graphics and Image Processing, August 19 1982.

19. Shneior, M., "Two Hierarchial Linear Feature Representations: Edge Pyramids and Edge Quadtrees", University of Maryland, TR 961, October 1980.

20. Zadeh, L. A., "A Theory of Approximate Reasoning", in Hayes, T. E., Mitchie, D., and Michulich, L. J. (EDS), Machine Intelligence, NY, John Wiley & Sons, Inc., 1979.

21. Duda, R. O., P. E. Hart, in B. L. Webber, N. J. Nilsson, "Subjective Bayesian Methods for rule-based inference systems", in Webber, B. L., N. J. (eds), Readings in Artificial Intelligence, 1981.

22. Garvey, T. D., J. D. Lawrence, and M. A. Fischer, "An inference technique for integrating knowledge from disparate sources", in Webber, B. L. Nilsson, N. J. (eds), Readings in Artificial Intelligence, 1981.

23. Cohen, P. R. and M. R. Grinberg, "A Theory of Heuristic Reasoning About Uncertainty," AI Magazine, Summer, 1983.

24. Department of the Army, Combat Inteligence, FM30-5, October 1973.

REFERENCES

(CONCLUDED)

25. SRI International, A Computer-Based Consultant for Mineral Exploration, SRI Project 6415, R. O Duda, P. E. Hart, K. Konolige, R. Reboh, September 1979.

26. Waterman, D. A. and F. Hayes-Roth, Pattern Directed Inference Systems, Academic Press, 1979.

26. R. Davis and D. B. Lenat, Knowledge-Based Systems in Artificial Intelligence, McGraw-Hill International, 1982.

27. Charnick, E., C. H. Riesbeck, D. V. McDermott, Artificial Intelligence Programming, Lawrence Erlbaun Associates, 1980.

28. Bonasso, R. P., "Running ANALYST on an APPLE", Working Notes, The MITRE Corporation, December 1982.

29. Weiss, S. M., and C. A. Kulikowski, "EXPERT: A System for Developing Consultation Models", Proceedings of the Sixth International Joint Conference on Artificial Intelligence, Tokoyo, 1979.

30. The MITRE Corporation, "The Role of Intelligence and All Source Analysis System in Battlefield Operations," B. R. Hart, MTR-81W00308, December 1981.

31. Department of the Army, Military Geographic Intelligence (Terrain), FM30-10.

32. Department of the Army, Special Applications of Terrain Intelligence, FM30-10A.

33. Department of the Army, Weather Support for Field Army Tactical Operations, FM31-3/AFM105-4.

34. Department of the Army, Opposing Forces Europe, FM30-102, November 1977.

35. Department of the Army, Intelligence Preparation of the Battlefield, 66000-ASupR, June 1980.

36. Department of the Army, Aggressor Order of the Battle, FM30-103.

DISTRIBUTION LIST

Internal

A-10 C. A. Zraket

D-10 K. E. McVicar

D-14 J. Dominitz
J. S. McManus
A. J. Tachmindji

D-73 C. Engelman

W-30 J. W. Benoit

W-70 P. G. Freck
R. P. Granato
R. A. Joy
F. W. Niedenfuhr
W. A. Tidwell

W-72 C. W. Sanders
R. D. Sexton
N. T. Hamisevicz

W-73 T. H. Nyman

W-74 H. J. Antonisse (10)
T. T. Bean
R. P. Bonasso (10)
D. R. Edmonds
M. Gale
B. R. Hart
C. R. Holt
E. P. Maimone
R. J. Nieporent
R. L. Hegerich

W-75 P. Messing

W70 Information Center

External

HQDA
ACSI DAMI-ZA
DAMI
DAMI-IS
DAMI-AML
DAMI-FRT (Mr. Beuch)
The Pentagon
Washington, DC 20301

HQDA
DUSA-OR (Mr. Hollis, Dr. Fallin)
DAMA-ZD (Mr. Woodall)
The Pentagon
Washington, DC 20301

Commander
U. S. Army Electronics Research and Development Command
2800 Powder Mill Road
Adelphi, MD 20783

Director, Electronic Warfare Lab
USAERADCOM
Fort Monmouth, NJ 07703

Director, Signals Warfare Lab
USAERADCOM
Vint Hill Farms Station
Warrenton, VA 22186

Defense Intelligence Agency
Room 1D918
The Pentagon
Washington, DC 20301

Jet Propulsion Laboratories
ATTN: Mr. S. Freesma
4800 Oak Grove
Pasadena, CA 91109

DISTRIBUTION LIST

(CONTINUED)

External (Continued)

RADC
ATTN: Dr. Northrop Fowler (COES)
ATTN: Andrew Kozak (IRDT)
Griffiss AFB, NY 13441

GRC
ATTN: Dennis Cooper
P.O. Box 6770
Santa Barbara, CA 93100

Stanford Research Institute
EK260
ATTN: Dr. T. Garvey
333 Ravenswood Avenue
Menlo Park, CA 94025

Mr. Ray Kirkwood
HQ TACATA
Ft. Hood, TX 76544
ATTN: ATCT-BA-SPS

The Army Model Management Office
U. S. Army Combined Arms Center
ATTN: ATZL-CAN-DO
COL Kenneth Wiersema (10)
ATTN: ATZL-CSC-I
Ft. Leavenworth, KS 66027

Army Library
ATTN: ANR-AL-RS
(Army Studies)
Room 1A518
The Pentagon
Washington, DC 20301

Commander
Defense Technical Information Center
ATTN: DDA
Cameron Station
Alexandria, VA 22314

Commandant
U. S. Army Command and General
Staff College
Ft. Leavenworth, KS 66027

Commandant
U. S. Army War College
Carlisle Barracks, PA 17013

Commander
U. S. Army Intelligence Center and School
ATTN: ATSL
(LTC Aikens, COL Fichtel, Carl Verhey,
MAJ Frew, CPT MacAdams, CPT Morgan,
CPT Garcia, CPT Weber)
Ft. Huachuca, AZ 85613

Commander
Soldier Support Center
ATTN: ATSG-DCD-AD
(LTC Crosby)
Ft. Benjamin Harrison, IN 46216

Director, U. S. Army Concepts
Analysis Agency
ATTN: CSCA-AZ (Mr. Hardison)
ATTN: CSCA-MC (Mr. Louer,
Mr. Shedlowsky, LTC Deems)
8120 Woodmont Avenue
Bethesda, MD 20014

Commander
Director, U. S. Army TRADOC
ATTN: ATAA-D (Mr. Goode)
ATAA-TG (Mr. Carrilo)
ATAA-TC (Mr. Matheson)
ATAA-TCF (Mr. McCoy)
White Sands, NM 88002

DISTRIBUTION LIST

(CONTINUED)

External (Continued)

Director, U. S. Army
Research Institute
ATTN: Dr. Ruth Phelps
Dr. Judy Englert
5001 Eisenhower Avenue
Alexandria, VA 22333

Bendix Commo Div.
1300 Joppa Road
Baltimore, MD 21204

Director, TRADOC Operations
Research Activity
White Sands, NM 88002

Deputy Commander Combined Arms
Operations Research Activity
ATTN: ATOR-CAA-DC (COL West)
ATOR-CAA-DR (Mr. Pleger)
ATOR-CAT-D (LTC Childs)
Ft. Leavenworth, KS 66027

Commander
Director, USA TRASANA
ATTN: ATORA-TCC (Mr. F. Payan,
Mr. A. Kientz)
White Sands, NM 88002

Commander
U. S. Army Intelligence and Threat
Analysis Center
ATTN: IAX-I (COL Satterwaite)
Arlington Hall Station, VA 22022

Commander
U. S. Army Training and
Doctrine Command
ATTN: ATCG-S (Mr. Christman)
ATCD-D (COL Shoffner)
Ft. Monroe, VA 23651

Applied Physics Lab
ATTN: Dr. V. Sigillito
Johns Hopkins Road
Laurel, MD 20707

ESL, Inc.
ATTN: Dr. Karmidjian
495 Java Drive
Sunnyvale, CA 94086

RCA
MS 10-7-6
ATTN: Jerry Sable
Camden, NJ 08102

Navy Center for Applied Research
in Artificial Intelligence
Naval Research Laboratory, Code 7510
ATTN: Ms. L. Davis, Dr. J. Slagle,
Dr. Y. T. Chien
4555 Overlook, S. W.
Washington, DC 20375

BDM Corporation
ATTN: Dr. J. Aldrich
One First National Plaza
Suite 1010
Dayton, Ohio 45402

Computer Sciences Department
University of Maryland
ATTN: Dr. L. Davis, Dr. A. Rosenfeld,
Dr. D. Nan, Dr. J. Reggia, Dr. L. Kanal
College Park, MD 20742

University Research Foundation
ATTN: Dr. S. Greenwood, Mr. G. Connor
6303 Ivy Lane, Suite 110
Greenbelt, MD 20770

DISTRIBUTION LIST

(CONCLUDED)

External (Concluded)

Director
Ballistic Research Laboratory
ATTN: DRDAR-BLB (Mort Hirshberg)
DRSMC-BLB (A) (Barry L. Reichand)
Aberdeen Proving Grounds
Aberdeen, MD 21005

Physics Department
University of Chicago
1118 E 58th Street
ATTN: Jim Wheeler
Chicago, ILL 60637

END

FILMED

02 — 84

DTIC